

Brain Gene Expression Analysis: a MATLAB toolbox for the analysis of brain-wide gene-expression data

Pascal Grange¹, Jason W. Bohland², Michael Hawrylycz³, Partha P. Mitra¹

¹ Cold Spring Harbor Laboratory, Cold Spring Harbor, New York 11724, United States

² Boston University, College of Health & Rehabilitation Sciences, Boston, MA 02215, United States

³ Allen Institute for Brain Science, Seattle, Washington 98103, United States

*E-mail: pascal.grange@polytechnique.org

November 28, 2012

Contents

| | |
|--|-----------|
| 1 The voxel-based Allen Atlas of the adult mouse brain | 5 |
| 1.1 Presentation of the dataset | 5 |
| 1.1.1 Co-registered <i>in situ</i> hybridization data | 5 |
| 1.1.2 Data matrices and gene filters | 5 |
| 1.2 From matrices to volumes | 6 |
| 1.2.1 The Allen Reference Atlas (ARA) | 6 |
| 1.2.2 From gene-expression vectors to volumes | 8 |
| 1.2.3 Visualization | 9 |
| 1.3 Relation between the various annotations in the ARA | 14 |
| 2 Genes versus neuroanatomy | 17 |
| 2.1 Localization scores | 17 |
| 2.1.1 Localization scores of a single gene in the ARA | 17 |
| 2.1.2 Localization scores of sets of genes in the ARA | 19 |
| 2.2 Fitting scores in the ARA | 20 |
| 2.2.1 Fitting scores of a single gene in the ARA | 20 |
| 2.2.2 Fitting scores of sets of genes in the ARA | 22 |
| 3 Genes versus genes: co-expression networks | 25 |
| 3.1 Co-expression network of genes in the Allen Atlas | 25 |
| 3.2 Monte Carlo analysis of brain-wide co-expression networks | 29 |
| 3.2.1 Special sets of genes versus full atlas | 29 |
| 3.2.2 Cumulative distribution function of co-expression coefficients in sets of genes drawn from the Allen Brain Atlas | 29 |
| 3.2.3 Thresholding the co-expression matrix | 31 |
| 3.2.4 Statistics of sizes of connected components | 32 |
| 3.2.5 Neuroanatomical properties of connected components | 33 |
| 4 Genes versus cell types (to be released in the second version of the toolbox, not publicly available as of November 2012) | 39 |
| 4.1 Cell-type-specific microarray data | 39 |
| 4.2 Brain-wide correlation profiles | 40 |
| 4.3 Brain-wide density estimates of cell types | 41 |

| | | |
|----------|--|-----------|
| 4.4 | Cell-type-based computational neuroanatomy | 42 |
| 4.5 | Visualization of correlation and density profiles of cell types | 44 |
| 5 | Clustering | 49 |
| 5.1 | Kullback–Leibler distance from gene-expression to a given probability distribution | 49 |
| 5.2 | Construction of a bipartite graph from a voxel-by-gene matrix | 49 |
| 6 | List of files | 55 |
| 6.1 | Preparation of data | 55 |
| 6.2 | Atlas navigation | 55 |
| 6.3 | Marker genes | 55 |
| 6.4 | Co-expression analysis | 56 |
| 6.5 | Clustering | 56 |
| 6.6 | Cell-type specific analysis | 56 |
| 6.7 | Miscellaneaous operations on matrices | 56 |
| 6.8 | Data files | 57 |
| 6.8.1 | Allen Brain Atlas, Allen Reference Atlas | 57 |
| 6.8.2 | Cell-type specific analysis | 57 |
| 6.9 | Saved results | 57 |
| 6.9.1 | Marker genes | 57 |
| 6.9.2 | Cell-type specific analysis | 57 |
| 6.9.3 | Playground | 58 |

Chapter 1

The voxel-based Allen Atlas of the adult mouse brain

1.1 Presentation of the dataset

1.1.1 Co-registered *in situ* hybridization data

The adult mouse brain is partitioned into $V = 49,742$ cubic voxels of side 200 microns, to which *in situ* hybridization data are registered [1, 2] for thousands of genes. For computational purposes, these gene-expression data can be arranged into a voxel-by-gene matrix.

For each voxel v , the *expression energy* of the gene g is a weighted sum of the greyscale-value intensities I evaluated at the pixels p intersecting the voxel:

$$E(v, g) = \frac{\sum_{p \in v} M(p)I(p)}{\sum_{p \in v} 1}, \quad (1.1)$$

where $M(p)$ is a Boolean mask that equals 1 if the gene is expressed at pixel p and 0 if it is not.

1.1.2 Data matrices and gene filters

The coronal atlas contains brain-wide data for $G_{\text{all}} = 4,104$ genes. The corresponding voxel-by-gene data matrix has size $V = 49,742$ by G_{all} , and is contained in the file ExpEnergy.mat. The list of genes arranged in the same order as the columns of the data matrix are obtained by using the function get_genes.m.

```
load('ExpEnergytop75percent.mat');
% the g-th column of voxel-by-gene matrix D corresponds
```

```

3% to the gene geneNamesAll( g )
4geneNamesAll = get_genes( Ref, 'top75CorrNoDup', 'allen' );
5% Entrez IDs are arranged in the same order
6% as gene names (unresolved Entrez ids are treated as zero)
7geneEntrezIdsAll = get_genes( Ref, 'top75CorrNoDup', 'entrez' );

```

The matrix in `ExpEnergyTop75Percent.mat` consists of the 3,041 columns of the matrix defined in Equation 1.1 that are best correlated with the corresponding genes in the sagittal atlas. The names and Entrez ids of the genes are obtained as follows:

```

1load( 'ExpEnergyTop75Percent.mat' );
2% the g-th column of voxel-by-gene matrix D
3% corresponds to the gene genesAllen( g )
4genesAllen = get_genes( Ref.Coral, 'top75corrNoDup', 'allen' );
5% Entrez IDs are arranged in the same order
6% as gene names (unresolved Entrez ids are treated as zero)
7genesEntrez = get_genes( Ref.Coral, 'top75corrNoDup', 'entrez' );

```

It should be noted that the Entrez ids that are not resolved are represented by zeroes in `geneEntrezIds`, so Entrez ids should be resolved *in fine* rather than used during computations.

The start-up file `mouse_start_up.m` loads the Allen Reference Atlas stored in the file `refAtlas.mat` and the data matrix `ExpEnergytop75percent.mat`, which can be used through the structure `Ref`.

- **Example 1. Systems of brain annotation.** The variable `cor=Ref.Coral` contains the coronal atlas, and it can be used to check some of the contents of Table 1.1. a detailed description of the fields in the structure `Ref.Coral.Annotations` can be found in the next section.

```

1mouse_start_up;
2cor = Ref.Coral;
3ann = cor.Annotations;
4display( ann );

```

1.2 From matrices to volumes

1.2.1 The Allen Reference Atlas (ARA)

The ARA comes in six different versions, described in Table 1.1. Each of these versions corresponds to an annotation of the three-dimensional grid by digital ids. Each of these digital ids corresponds to a brain region. The correspondence between ids and names of brain regions can be resolved using the fields `Ref.Coral.ids` and `Ref.Coral.labels`.

| Identifier index | Identifier (name of annotation) | Number of regions | Hierarchical? | Number of annotated voxels |
|------------------|---------------------------------|-------------------|---------------|----------------------------|
| 1 | standard | 209 | Yes | 49,742 |
| 2 | cortex | 40 | Yes | 11,862 |
| 3 | standard+cortex | 242 | Yes | 49,742 |
| 4 | fine | 94 | No | 22,882 |
| 5 | big12 | 13 | No | 25,155 |
| 6 | cortexLayers | 8 | No | 11,862 |

Table 1.1: Systems of annotations of the adult mouse brain in the digital version of the Allen Reference Atlas, at a resolution of 200 microns.

- **Example 2. From brain regions to digital ids in the ARA.** Consider the `fine` annotation (identifier index 4 in Table 1) and let us work out a three-dimensional grid with ones at the voxels corresponding to the caudoputamen in this annotation, and zeroes everywhere else. We can use this grid to compute the number of voxels in the caudoputamen.

```

1% the three-dimensional grid containing the fine annotation of the brain
2annotFine = get_annotation( cor, 'fine' );
3% the names of brain regions
4labels = ann.labels{ 4 };
5% the numerical ids of the regions
6ids = ann.ids{ 4 };
7% where is the caudoputamen in the list? There may be spaces
8% in the labels, skip them
9labelsNoSpace = regexp替换( labels, '\W', '' );
10caudoputamenIndex = find( strcmp( labelsNoSpace, 'Caudoputamen' ) == 1 );
11caudoputamenId = ids( caudoputamenIndex );
12% put zeros at all the voxels that do not belong to caudoputamen
13volCaudoputamen = annotFine;
14volCaudoputamen( volCaudoputamen ~= caudoputamenId ) = 0;
15volCaudoputamen( volCaudoputamen == caudoputamenId ) = 1;
16% count the voxels in caudoputamen
17numVoxInCaudoputamen = sum( sum( sum( volCaudoputamen ) ) );
18display( numVoxInCaudoputamen )

```

The value of the variable `numVoxInCaudoputamen` computed in the above code snippet should be 1248.

1.2.2 From gene-expression vectors to volumes

At a resolution of 200 microns, the Allen Reference Atlas (ARA) is embedded in a three-dimensional grid of size $67 \times 41 \times 58$. Out of the $V_{\text{tot}} = 67 \times 41 \times 58$ in the grid, $V = 49,742$ are in in the brain according to the ARA.

The brain voxels can be mapped to a three-dimensional grid using the function `make_volume_from_labels.m` and a specified voxel filter, corresponding to one of the versions of the ARA (as per Table 1.1).

- **Example 3. Whole-brain filter.** Let us take the whole first column of the data matrix and map it to a three-dimensional grid, using the whole-brain filter, corresponding to the 'standard' annotation:

```

1wholeBrainFilter = Ref.Coronal.Annotations.Filter{ 1 };
2display( wholeBrainFilter );
3brainFilter = get_voxel_filter( cor, wholeBrainFilter );
% a column vector with 49,742 elements
5col1 = D( :, 1 );
6display( size( col1 ) )
% map this column vector to a volume
8vol1 = make_volume_from_labels( col1, brainFilter );
9display( size( vol1 ) );

```

Some of the annotations do not extend to the whole brain, as can be seen from Table 1.1. when comparing columns of the matrix of gene-expression energies to regions in the ARA, it is important to restrict the matrix to the rows that correspond to annotated voxels. For each system of annotation, the field `Ref.Coronal.Annotations.Filter` is the list of voxels in a $67 \times 41 \times 58$ grid (*not* a list of row indices in the matrix of expression energies), that are annotated. The example below shows how to recover the list of row indices in the gene-expression matrix corresponding to a given filter.

- **Example 4. More filters.** For instance, one can check that the following two snippets produce the same matrix `DFiltered`, consisting of the rows of the full gene-expression matrix corresponding to voxels in the `fine` annotation:

```

1 % Compute the set of rows using a volume of
2 %indices and the filter
3 cor = Ref.Coronal;
4 identifierIndex = 4;
5 brainFilter = get_voxel_filter( cor, 'brainVox' );
6 numVox = numel( brainFilter );
7 % label the voxels by integers
8 indsBrainVoxels = 1 : numVox;
9 % arrange the integers in a volume

```

```

10 indsVol = make_volume_from_labels( indsBrainVoxels, brainFilter );
11 filter = get_voxel_filter( cor, ann.filter{ identifierIndex } );
12 % restrict the volume to the voxels that are in the filter
13 indsFiltered = indsVol( filter );
14 DFiltered = D( indsFiltered, : );}

```

```

1 %apply the filter to each column of the data matrix
2 cor = Ref.Coronal;
3 identifierIndex = 4;
4 brainFilter = get_voxel_filter( cor, 'brainVox' );
5 filter = get_voxel_filter( cor, ann.filter{ identifierIndex } );
6 numGenes = size( D, 2 );
7 % restrict each column of the data matrix to voxels
8 % that are in the filter
9 for gg = 1 : numGenes
10     geneVol = make_volume_from_labels( D( :, gg ), brainFilter );
11     DFiltered( :, gg ) = geneVol( filter );
12 end

```

Note also that the voxel filter can be recomputed from the three-dimensional annotation and the brain-wide filter:

```

1filterStored = get_voxel_filter( cor, 'leftVox' );
2brainFilter = get_voxel_filter( cor, 'brainVox' );
3annot = get_annotation( cor, ann.identifier{ identifierIndex } );
4annotArray = annot( brainFilter );
5% the annotated voxels are filled with non-zero numerical ids
6annotatedIndices = find( annotArray ~= 0 );
7filterComputed = brainFilter( annotatedIndices );

```

1.2.3 Visualization

Maximal-intensity projection

Given a three-dimensional grid of size $67 \times 41 \times 58$, its maximal-intensity projection onto sagittal, coronal and axial planes can be plotted using the function

`plot_intensity_projections.m`

- **Example 5. Maximal-intensity projections of the expression energy of a gene.** Let us plot maximal-intensity projections of the expression energy of *Gabra6*. The abovefollowing code snippet should reproduce Figure 1.1.

```

1% find out to which column of the data matrix it corresponds
2indexGabra6 = find( strcmp( genesAllen, 'Gabra6' ) == 1 );
3columnGabra6 = D( :, indexGabra6 );
4% map the column to a volume

```

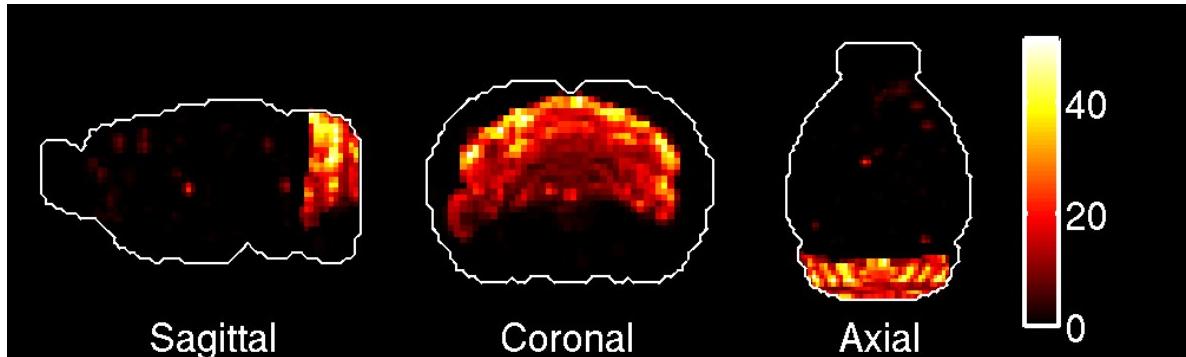


Figure 1.1: Maximal-intensity projections of the expression energy of *Gabra6*.

```
5 volGabra6 = make_volume_from_labels( columnGabra6, brainFilter );
6 plot_intensity_projections( volGabra6 );
```

- **Example 6. Maximal-intensity projections of brain regions** We can use the function `plot_intensity_projections.m` to visualize brain regions in the atlas. For instance, we can use the grid `volCaudoputamen` computed in one of the examples above and reproduce Figure 1.2.

```
1 % the three-dimensional grid containing the fine annotation of the brain
2 annotFine = get_annotation( cor, 'fine' );
3 labels = ann.labels{ 4 };
4 ids = ann.ids{ 4 };
5 labelsNoSpace = regexp替换( labels, '\W', '' );
6 caudoputamenIndex = find( strcmp( labelsNoSpace, 'Caudoputamen' ) == 1 );
7 caudoputamenId = ids( caudoputamenIndex );
8 volCaudoputamen = annotFine;
9 volCaudoputamen( volCaudoputamen ~= caudoputamenId ) = 0;
10 volCaudoputamen( volCaudoputamen == caudoputamenId ) = 1;
11 plot_intensity_projections( volCaudoputamen );
```

Example 7. The `fineAnatomy` filter. It is manifest from the projection of the characteristic function of the caudoputamen in the `fine` annotation (Figure 1.2) that the `fine` annotation contains only the left hemisphere of the brain. We can confirm this by applying the voxel filter of the `fine` annotation (called `fineAnatomy`, which is the value of `Ref.Coronal.Annotations.filter{4}`) to the gene-expression vector of *Gabra6*. The following snippet should reproduce Figure 1.3, which is equivalent to Figure 4.1, with all the voxels that are not in the `fineAnatomy` filter filled with zeros.

```
1 % take the brain-wide expression of Gabra6
2 indexGabra6 = find( strcmp( genesAllen, 'Gabra6' ) == 1 );
3 columnGabra6 = D( :, indexGabra6 );
4 volGabra6 = make_volume_from_labels( columnGabra6, brainFilter );
5 % consider the voxel filter of the 'fine' annotation
```

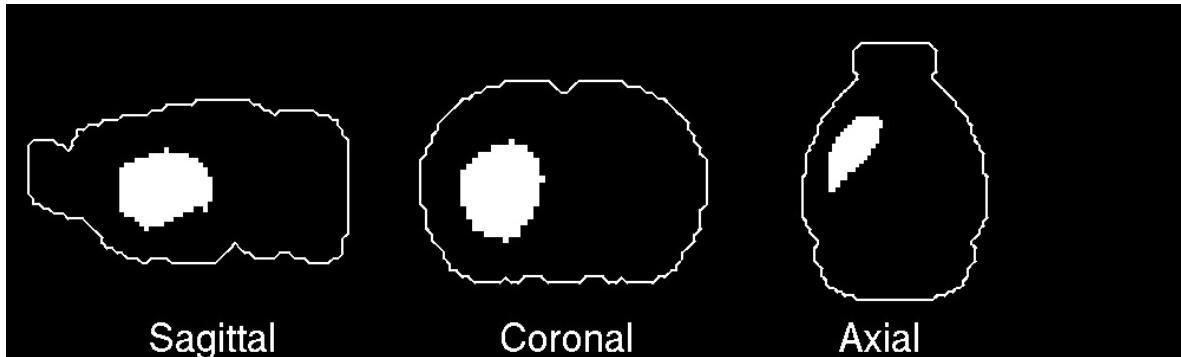


Figure 1.2: Maximal-intensity projections of the characteristic function of the caudoputamen in the `fine` annotation.

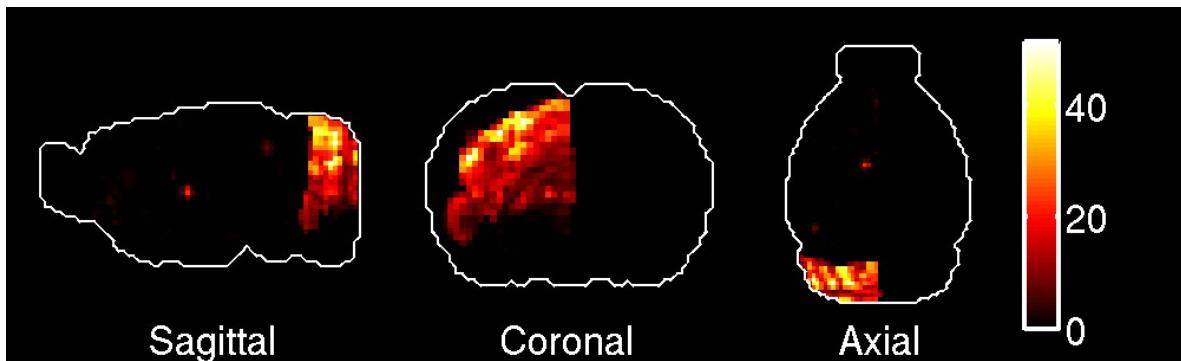


Figure 1.3: Maximal-intensity projections of the expression energy of *Gabra6*.

```

6fineFilter = get_voxel_filter( cor, 'fineAnatomy' );
% take the expression values at the voxels in the filter , and arrange them
% in a column vector (one can check that
% dataGabra6Filtered has the same number of
10% elements as fineFilter
11colGabra6Filtered = volGabra6( fineFilter );
12% map these values to a three-dimensional grid , using fineFilter
13volGabra6Filtered = make_volume_from_labels( colGabra6Filtered, ...
14 fineFilter );
15plot_intensity_projections( volGabra6Filtered );

```

- **Example 8. Brain-wide versus left hemisphere.** As can be seen on Table 1.1, the standard annotation contains all the voxels in the brain (49,742 of them). We can compute the characteristic function of the caudoputamen in this annotation, and check that the left half of the caudoputamen coincides with Figure 1.2, and that the number of voxels is twice the number computed in the `fine` annotation. The code snipped below should reproduce Figure 1.4.

```
1identifierIndex = 1;
```

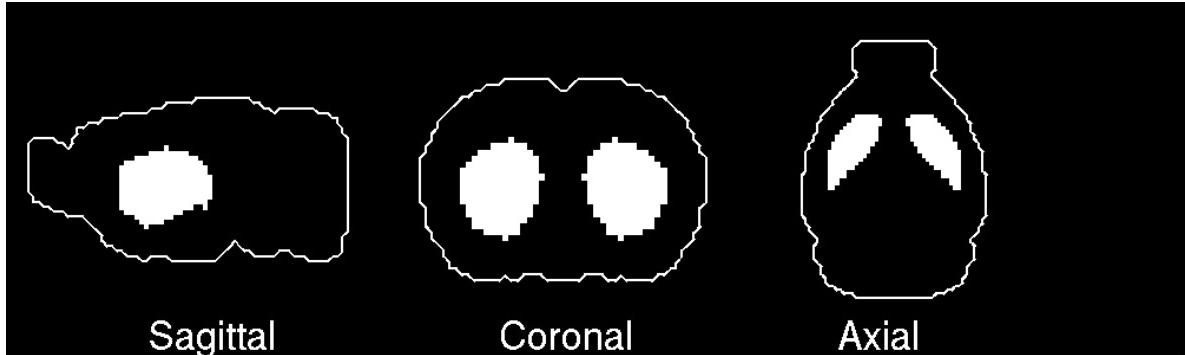


Figure 1.4: Maximal-intensity projections of the caudoputamen in the standard annotation.

```

2annotStandard = get_annotation( cor, ann.identifier{ identifierIndex } );
3ann = cor.Annotations;
4labels = ann.labels{ identifierIndex };
% their numerical ids
5ids = ann.ids{ identifierIndex };
% where is the caudoputamen in the list?
6labelsNoSpace = regexp替換( labels, '\W', '' );
7caudoputamenIndex = find( strcmp( labelsNoSpace, 'Caudoputamen' ) == 1 );
8caudoputamenId = ids( caudoputamenIndex );
% put zeros at all the voxels that do not belong to caudoputamen
9volCaudoputamen = annotStandard;
10volCaudoputamen( volCaudoputamen ~= caudoputamenId ) = 0;
11volCaudoputamen( volCaudoputamen == caudoputamenId ) = 1;
% count the voxels in caudoputamen
12numVoxInCaudoputamenStandard = sum( sum( sum( volCaudoputamen ) ) );
13display( numVoxInCaudoputamenStandard );
% reproduce Figure \ref{projectionCaudoputamenStandard}
14plot_intensity_projections( volCaudoputamen );

```

Sections

The function `flip_through_sections.m` allows to go through the sections of a $(67 \times 41 \times 58)$ volume, of a kind specified in the options. It pauses between sections. The duration of the pause is one second by default, it can be adjusted using the field `secondsOfPause` of the options. If the value of `secondsOfPause` is negative, the user will have to press a key to display the next section.

- **Example 9.** Sections of the average across all genes in the dataset. The following code allows to visualize the coronal, sagittal and axial sections of the average expression

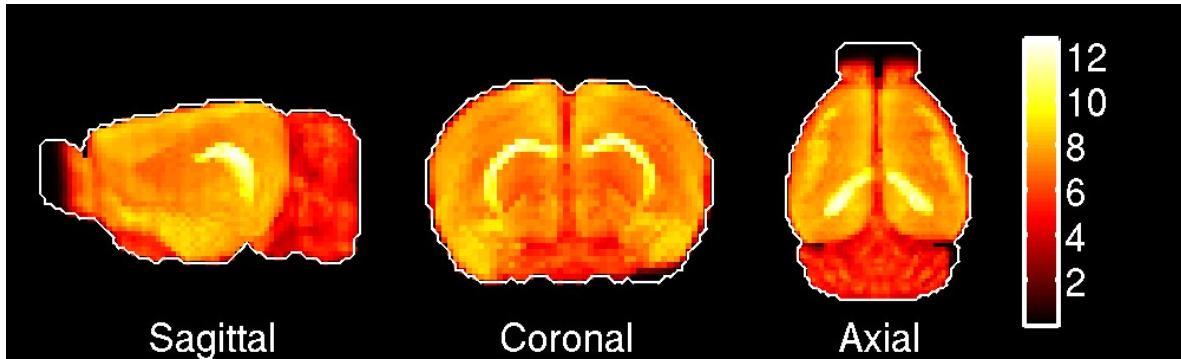


Figure 1.5: Maximal-intensity projections of the average gene-expression \bar{E} defined in Equation 1.2.

\bar{E} across all genes in the data matrix, defined in Equation 1.2:

$$\bar{E}(v) = \frac{1}{G} \sum_{g=1}^G E(v, g). \quad (1.2)$$

It also reproduces the maximal-intensity projections of \bar{E} , shown in Figure 1.5.

```

1% the average gene-expression vector;
2numGenes = size( D, 2 );
3avgExpressionVector = sum( D, 2 ) / numGenes;
4% map it to a volume
5brainFilter = get_voxel_filter( Ref.Coral, 'brainVox' );
6avgExpressionVol = make_volume_from_labels( avgExpressionVector, brainFilter );
7% visualize projections of this volume
8splot_intensity_projections( avgExpressionVol );
9hold off;
10%flip through coronal sections, pause one second between sections
11optionsCoronal = struct( 'sectionStyle', 'coronal', 'secondsOfPause', 1 );
12flipThroughSections = flip_through_sections( avgExpressionVol, optionsCoronal );
13close all;
14%flip through sagittal sections
15optionsSagittal = struct( 'sectionStyle', 'sagittal', 'secondsOfPause', 1 );
16flipThroughSections = flip_through_sections( avgExpressionVol, optionsSagittal );
17close all;
18%flip through axial sections
19optionsAxial = struct( 'sectionStyle', 'axial', 'secondsOfPause', 1 );
20flipThroughSections = flip_through_sections( avgExpressionVol, optionsAxial );
21close all;

```

1.3 Relation between the various annotations in the ARA

The field `Ref.Coronal.Annotations.parentIdx` in the reference data structure gives the indices (not the numerical id) of the parents of the regions in a given annotation, arranged in the same order as these regions. If a region does not have a parent in the considered system of annotation, the corresponding entry in `parentIdx` is zero.

NB: the list of regions in the 'cortex' annotation is not closed under the operation of taking the parent of a region in a hierarchy. The field `parentSymbols` has to be used instead of `parentIdx` to work out the hierarchy.

- **Example 10. Hierarchical and non-hierarchical annotations.** Let us check that the `big12` and `fine` annotations are non-hierarchical, and investigate parents and descendants of the caudoputamen and of the cerebellum in the `standard` annotation.

```

1ann = Ref.Coronal.Annotations;
2% the parent indices in the 'big12' annotation are zeroes
3parentIdxBig12 = ann.parentIdx{ 5 }
4% check that the 'big12' annotation is non-hierarchical
5Big12IsNonHierarchical = isempty( setdiff( unique( parentIdxBig12 ),...
6 [ 0 ] ) )
7% same goes for the 'fine' annotation
8parentIdxFine = ann.parentIdx{ 4 }
9% check that the 'fine' annotation is non-hierarchical
10fineIsNonHierarchical = isempty( setdiff( unique( parentIdxBig12 ),...
11 [ 0 ] ) )
12% work out the names of the parents of a region in the standard
13% annotation
14labelsStandard = ann.labels{ 1 };
15labelsStandardNoSpace = regexp替換( labelsStandard, '\W', '' );
16indexCaudoputamen = find( strcmp( labelsStandardNoSpace, ...
17 'Caudoputamen' ) == 1 );
18parentIndicesStandard = ann.parentIdx{ 1 };
19parentIndexCaudoputamen = parentIndicesStandard( indexCaudoputamen );
20display( labelsStandard( parentIndexCaudoputamen ) )
21% and the names of the descendants
22caudoputamenDescendantIndices =
23find( parentIndicesStandard == indexCaudoputamen );
24% Is caudoputamen a leaf in the tree underlying the annotation?
25caudoputamenIsALeaf = isempty( caudoputamenDescendantIndices );
26display( caudoputamenIsALeaf );
27% What are the subregions of the cerebellum?
28cerebellumIndex = find( strcmp( labelsStandardNoSpace, 'Cerebellum' ) == 1 );
29cerebellumDescendantIndices = find( parentIndicesStandard == cerebellumIndex );
30cerebellumDescendants = labels( cerebellumDescendantIndices );
31display( cerebellumDescendants );

```

- **The fine annotation compared to the big12 annotation.** Let us show that

the '`fine`' annotation is a refinement of the `big12` annotation. It is manifest from Table 1.1 that the `fine` annotation covers fewer voxels than, but we can check that 1) all these voxels are also in the `big12` annotation, and 2) each region in the `fine` annotation intersects only one region in the `big12` annotation.

```

% take the fine and big annotation, their ids and labels
2 identifierIndex = 4
3 cor = Ref.Coral;
4 ann = cor.Annotations;
5 identifierIndexBig = 5;
6 annotBig = get_annotation( cor, ann.identifier{ identifierIndexBig } );
7 idsBig = ann.ids{ identifierIndexBig };
8 labelsBig = ann.labels{ identifierIndexBig };
9 annotFine = get_annotation( cor, ann.identifier{ identifierIndex } );
10 idsFine = ann.ids{ identifierIndex };
11 labelsFine = ann.labels{ identifierIndex };
12 % how are the voxels in a given fine annotation labelled in big12
13 for ff = 1 : numel( idsFine )
14     idFine = idsFine( ff );
15     annotLoc = annotFine;
16     annotLoc( annotLoc ~= idFine ) = 0;
17     valsBig = unique( annotBig( find( annotFine == idFine ) ) );
18     % if this region is not included in one of the big12 region,
19     % valsBig contains 0 and/or several integers, and the following
20     % causes an error
21     if numel( valsBig ) ~= 1 && ~isempty( find( valsBig == 0 ) )
22         errorAtlasToBig = 'The annotation is not a refinement of big12';
23     else
24         indexParentBigAtlas( ff ) = find( idsBig == valsBig );
25     end
26 end
27 labelsParentBigAtlas = labelsBig( indexParentBigAtlas );
28 annotationFineToBig12.labelsFine = labelsFine;
29 annotationFineToBig12.indexParentBigAtlas = indexParentBigAtlas;
30 annotationFineToBig12.labelsParentBigAtlas = labelsParentBigAtlas;

```

Now from `big12` to `fine`:

```

% take the fine and big annotation, their ids and labels
2 identifierIndex = 4
3 cor = Ref.Coral;
4 ann = cor.Annotations;
5 identifierIndexBig = 5;
6 annotBig = get_annotation( cor, ann.identifier{ identifierIndexBig } );
7 idsBig = ann.ids{ identifierIndexBig };
8 labelsBig = ann.labels{ identifierIndexBig };
9
10 annotFine = get_annotation( cor, ann.identifier{ identifierIndex } );
11 idsFine = ann.ids{ identifierIndex };
12 labelsFine = ann.labels{ identifierIndex };

```

16CHAPTER 1. THE VOXEL-BASED ALLEN ATLAS OF THE ADULT MOUSE BRAIN

```

13for bb = 1 : numel( idsBig )
14    idBig = idsBig( bb );
15    valsFine = unique( annotFine( find( annotBig == idBig ) ) );
16    indicesFine = [];
17    for vv = 1 : numel( valsFine )
18        valFine = valsFine( vv );
19        if valFine ~= 0
20            indicesFine = [ indicesFine , find( idsFine == valFine ) ];
21        end
22    end
23    annotationBig12ToFine.indicesInFineAtlas{ bb } = indicesFine;
24    annotationBig12ToFine.subregionsInFineAtlas{ bb } = labelsFine( indicesFine );
25    annotationBig12ToFine.labelBig{ bb } = labelsBig{ bb };
26end

```

The toolbox contains these code snippets as two functions that work out the refinement of the `big12` annotation, and the organisation of the `fine` annotation into larger regions of the brain. One can check that the following reproduces the above results:

```

1annotationFineToBig12 = annotation_fine_to_big12( Ref );
2annotationBig12ToFine = annotation_big12_to_fine( Ref );

```

For instance, the 16-th region in the `fine` is Nucleus Accumbens. Check it:

```

1display( Ref.Coronal.Annotations.labels{ 4 }( 16 ) );
2% check that this region is a the right index in annotationFineToBig12
3display( annotationFineToBig12.labelsFine( nAccIndex ) );
4% In which region of Big12 is it included?
5display( annotationFineToBig12.labelsParentBigAtlas( nAccIndex ) );
6% Apart from nucleus accumbens which are the other subregions of the
7% striatum?
8striatumIndexInBig12 = annotationFineToBig12.indexParentBigAtlas( nAccIndex );
9% check that the striatum is a the right index in annotationBig12ToFine
10display( annotationBig12ToFine.labelBig{ striatumIndexInBig12 } );
11%the subregions of the striatum are the following
12display( annotationBig12ToFine.subregionsInFineAtlas{ striatumIndexInBig12 } );

```

Chapter 2

Genes versus neuroanatomy

2.1 Localization scores

2.1.1 Localization scores of a single gene in the ARA

Let us define [3] the localization score $\lambda_\omega(g)$ of a gene g in a region ω as the ratio of the squared L^2 -norm of the expression energy of gene g in region ω to the squared L^2 -norm of the expression energy of gene g in the set Ω of voxels that are annotated in the version of the ARA that contains ω .

$$\lambda_\omega(g) = \frac{\sum_{v \in \omega} E(v, g)^2}{\sum_{v \in \Omega} E(v, g)^2}, \quad (2.1)$$

It can be computed from a voxel-by-gene matrix (with one score per column for a fixed region ω), using the function `localization_from_id.m`, given the numerical id of the region ω and the numerical identifier corresponding to an annotation containing ω .

- **Example 11. One gene, one region.** Let us compute the localization score of *Pak7* in the cerebral cortex, as defined by the `big12` annotation of the left hemisphere.

```
1cor = Ref.Coronal;
2ann = cor.Annotations;
3identifierIndex = 5;
4ids = ann.ids{ identifierIndex };
5labels = ann.labels{ identifierIndex };
% where is the Cerebral cortex in the atlas?
7cortexIndex = find( strcmp( labels, 'Cerebral cortex' ) );
8idCortex = ids( cortexIndex );
9indexPak7 = find( strcmp( genesAllen, 'Pak7' ) == 1 );
10localizationPak7Cortex = localization_from_id( Ref, ...
11 D( :, indexPak7 ), identifierIndex, idCortex );
```

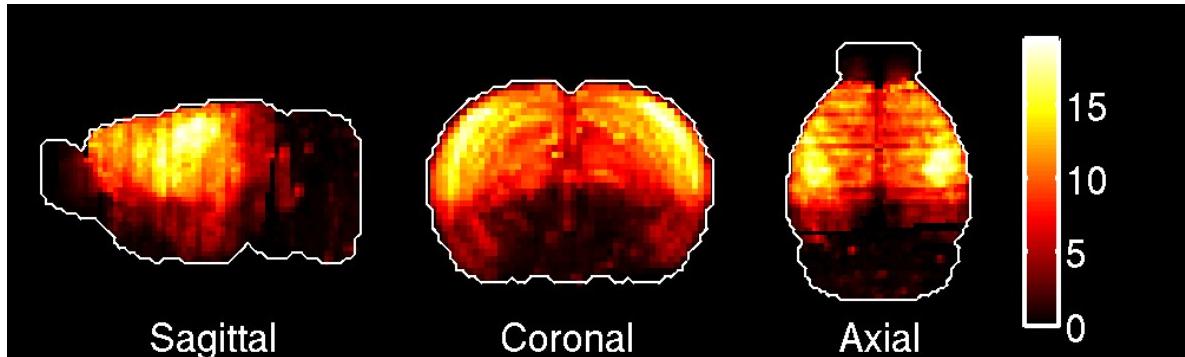


Figure 2.1: Maximal-intensity projection of the best-localized gene in the cerebral cortex.

The function `localization_scores_region_by_gene.m` computes the localization scores of all the columns of the data matrix in a given system of annotation. The set of voxels called `Basic cell groups` and `regions` can be skipped by choosing `struct('indInit', 2)` as the last argument of the function

`localization_scores_region_by_gene.m` in the '`big12`' annotation. The use of this function is illustrated .

- **Example 13. More genes, more regions: comparison to the atlas by localization scores.** The code snippet below should reproduce Figure 2.1 among other things.

```

1 identifierIndex = 5;
2 localizationRegionByGeneBig12 = localization_scores_region_by_gene( Ref ,...
3 D, identifierIndex, struct( 'indInit', 1 ) );
4 % check that the columns of the matrix of localization
5 % scores sum to 1
6 numGenes = size( locScores, 2 );
7 shouldBeOne = mean( sum( locScores ) );
8 display( shouldBeOne );
9 % look at the best-localized gene in each brain region
10 for rr = 2 : numel( ids )
11     labelReg = labels( rr );
12     display( labelReg );
13     regionScores = locScores( rr, : );
14     bestGeneIndex = find( regionScores == max( regionScores ) );
15     bestGeneIndex = bestGeneIndex( 1 );
16     display( genesAllen( bestGeneIndex ) );
17     bestExpressionByLocalization = D( :, bestGeneIndex );
18     bestExpressionByLocalization =...
19         make_volume_from_labels( bestExpressionByLocalization, brainFilter );
20     plot_intensity_projections( bestExpressionByLocalization );
21     saveas( gcf, [ 'singleLocMarkerIdentifier',...
22             num2str( identifierIndex ), 'Region', num2str( rr ) ], 'png' );
23     pause( 1 );
24 end

```

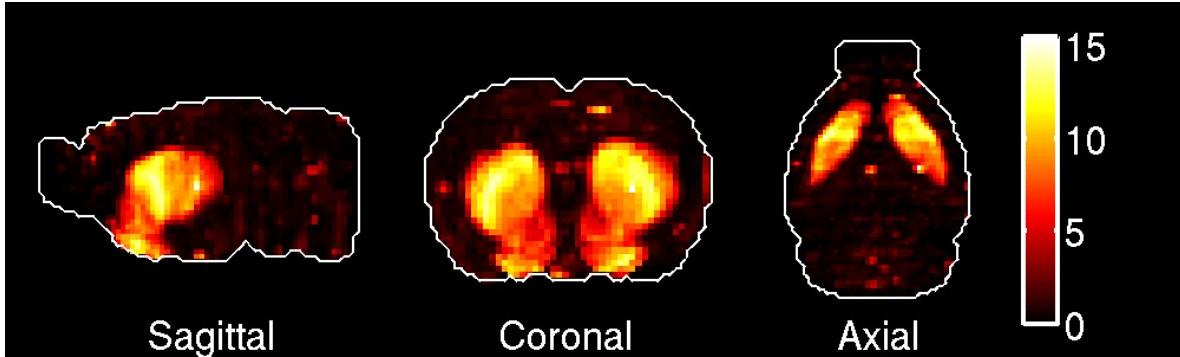


Figure 2.2: Maximal-intensity projection of the best-localized gene in the striatum.

2.1.2 Localization scores of sets of genes in the ARA

Equation 2.1 is naturally generalised to a linear combination of gene-expression vectors, weighted by real coefficients:

$$E_\alpha(v) := \sum_{g=1}^G \alpha_g E(v, g), \quad \alpha = (\alpha_1, \dots, \alpha_G) \in \mathbf{R}^G, \quad (2.2)$$

where G is the number of genes in our dataset ($G = 3,041$ by default when using the start-up file `mouse_start_up.m`).

Let us define the localization score in the brain region ω of a weighted set of genes encoded by Equation 2.2 as

$$\lambda_\omega(\alpha) = \frac{\sum_{v \in \omega} \left(\sum_g \alpha_g E(v, g) \right)^2}{\sum_{v \in \Omega} \left(\sum_g \alpha_g E(v, g) \right)^2} = \frac{\alpha^t J^\omega \alpha}{\alpha^t J^\Omega \alpha}, \quad (2.3)$$

- **Example 14.** Let us compute the best generalized localization scores in the `big12` annotation.

```
% Consider the big12 annotation
2 identifierIndex = 5;
3 optionsGen = struct( 'recomputeQuad', 0, 'saveResults', 0, 'recomputeQuadTot', 0 );
% Cortex only
5 for rr = 2
% Whole atlas
% for rr = 1 : numel( Ref.Coral.Annotations.ids{ identifierIndex } )
8     localizationGeneralizedRegionBig12{ rr } = localization_generalized( Ref, ...
9         D, identifierIndex, rr, optionsGen );
10    genEigenVector = localizationGeneralizedRegionBig12{ rr }.genVec;
11    weightedSumOfGenes = D * genEigenVector;
12    vol = make_volume_from_labels( weightedSumOfGenes, brainFilter );
```

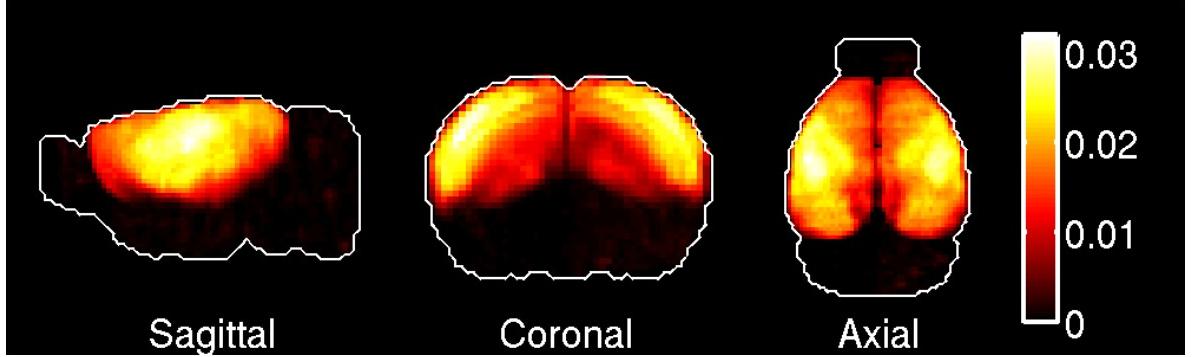


Figure 2.3: Maximal-intensity projection of the best generalized marker if the cerebral cortex.

```

13 %plot the expression of the solution
14 plot_intensity_projections( vol );
15 saveas( gcf, [ 'genLocMarkerIdentifier', num2str( identifierIndex ),...
16 'Region', num2str( rr ) ], 'png' );
17 pause( 3 );
18 end

```

2.2 Fitting scores in the ARA

2.2.1 Fitting scores of a single gene in the ARA

The fitting score $\phi_\omega(g)$ of a gene g in a region ω is defined

$$\phi_\omega(g) = 1 - \frac{1}{2} \sum_{v \in \Omega} (E_g^{\text{norm}}(v) - \chi_\omega(v))^2, \quad (2.4)$$

where E_g^{norm} is the L^2 -normalized g -th column E_g of the matrix of gene-expression energies:

$$E_g^{\text{norm}}(v) = \frac{E(v, g)}{\sqrt{\sum_{w=1}^V E(w, g)^2}}, \quad (2.5)$$

the symbol Ω denotes the set of voxels in a given system of annotation that contains region ω . The definition is the only decreasing affine function of the squared L^2 norm of the difference between the normalized gene-expression vector of gene g and the characteristic function χ_ω of the region ω :

$$\chi_\omega(v) = \frac{\mathbf{1}(v \in \omega)}{\sqrt{\sum_{w \in \omega} \mathbf{1}(w \in \omega)}}, \quad (2.6)$$

where the denominator in Equation 2.6 ensure the L^2 -normalization $\sum_{v=1}^V \chi_\omega(v)^2 = 1$.

Given a gene, a system of annotation chosen among the ones in Table 1.1, and the numerical id of a region in this system of annotation, the function `fitting_from_id.m` computes the fitting score of the gene to the corresponding region.

- **Example 15. One gene, one region in a given annotation.** Let us compute the fitting score of *Pak7* in $\omega = \text{Cerebral cortex}$, in the `big12` annotation (hence `identifierIndex = 5`, see Table 1.1).

```

1 ann = cor.Annotations;
2 identifierIndex = 5;
3 ids = ann.ids{ identifierIndex };
4 labels = ann.labels{ identifierIndex };
5 % where is the Cerebral cortex in the atlas?
6 cortexIndex = find( strcmp( labels, 'Cerebral cortex' ) );
7 idCortex = ids( cortexIndex );
8 indexPak7 = find( strcmp( genesAllen, 'Pak7' ) == 1 );
9 fittingPak7Cortex = fitting_from_id( Ref, D( :, indexPak7 ),...
10 identifierIndex, idCortex );

```

The function `fitting_from_id.m` can also used to compute the fitting score of several genes. If the second argument of the function is a voxel-by-gene matrix with p columns, the function returns an array of p fitting scores arranged in the same order as the columns.

Given a version of the ARA (specified by the index `identifierIndex`), a region-by-gene matrix of all the fitting scores of all genes corresponding to the columns of the data matrix. Note that the columns of this region-by-gene score matrix do not sum to a constant (the squares of the entries of each column sum to the square of the fraction of the gene-expression that projects onto the set of voxels in the annotation, which is at most 1).

- **Example 16. Fitting scores of all genes in all the regions in the `big12` annotation.**

```

1 identifierIndex = 5;
2 ids = ann.ids{ identifierIndex };
3 labels = ann.labels{ identifierIndex };
4 fittingScoresRegionByGene = fitting_scores_region_by_gene( Ref, ...
5 D, identifierIndex, options );
6 fittingScores = fittingScoresRegionByGene.fittingScores;
7 for rr = 2 : numel( ids )
8     labelReg = labels( rr );
9     display( labelReg );
10    fittingScoresRegionByGene = fittingScoresRegionByGeneBig12.fittingScores;
11    regionScores = fittingScoresRegionByGene( rr, : );
12    bestGeneIndex = find( regionScores == max( regionScores ) );
13    bestGeneIndex = bestGeneIndex( 1 );
14    display( genesAllen( bestGeneIndex ) );

```

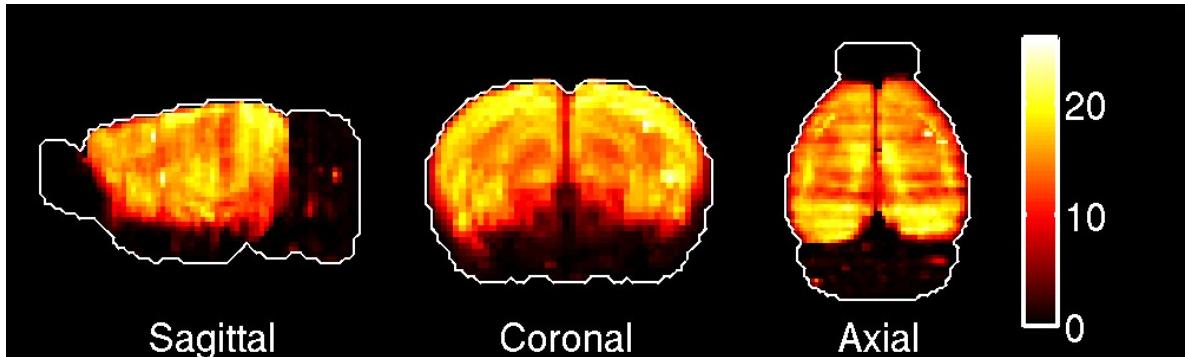


Figure 2.4: Maximal-intensity projection of the best-fitted gene in the cerebral cortex.

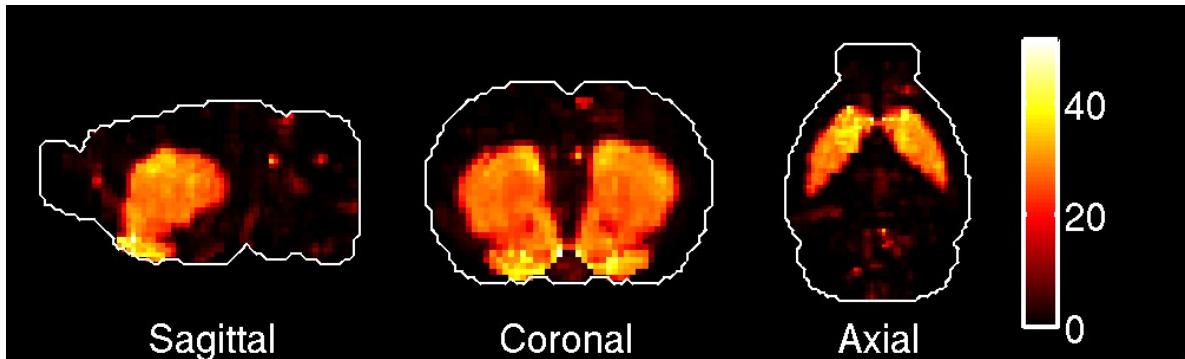


Figure 2.5: Maximal-intensity projection of the best-fitted gene in the striatum.

```

15 bestExpressionByFitting = D( :, bestGeneIndex );
16 bestExpressionByFitting = make_volume_from_labels( bestExpressionByFitting, ...
17   brainFilter );
18 plot_intensity_projections( bestExpressionByFitting );
19 pause( 1 );
20 end

```

2.2.2 Fitting scores of sets of genes in the ARA

Like the localization score, the fitting score can be generalized to linear combinations of sets of genes:

$$\phi_\omega(\alpha) = 1 - \frac{1}{2} \sum_{v \in \Omega} (E_\alpha^{\text{norm}}(v) - \chi_\omega(v))^2, \quad \alpha = (\alpha_1, \dots, \alpha_G) \in \mathbf{R}_+^G, \quad (2.7)$$

where E_α^{norm} is the L^2 -normalized gene-expression vector corresponding to the coefficients $(\alpha_1, \dots, \alpha_G)$

$$E_\alpha^{\text{norm}}(v) = \frac{E_\alpha(v)}{\sqrt{\sum_{v \in \Omega} E_\alpha(v)^2}}. \quad (2.8)$$

$$\text{ErrFit}_{L^1-L^2}^{\omega, \Lambda}(\{\alpha\}) = \|E_{\alpha}^{\text{norm}} - \chi_{\omega}\|_{L^2}^2 + \Lambda \|\alpha\|_{L^1}, \quad (2.9)$$

which can be minimized wrt the weights of the genes using Matlab code implementing an interior-point method by Koh [4]:

$$\alpha_{\omega}^{\Lambda} = \operatorname{argmin}_{\alpha \in \mathbf{R}_+^G} \text{ErrFit}_{L^1-L^2}^{\omega, \Lambda}(\{\alpha\}). \quad (2.10)$$

- **Example 17. Best-fitted sets of genes in (some regions of) the big12 annotation.** The code below should reproduce Figures ?? and ??

```

1 optionsSet = struct( 'lambdaVals', [ 0.005 ], ...
2     'identifierIndex', 5, 'numGenes', 20 );
3 lambdaVals = optionsSet.lambdaVals;
4 numGenes = optionsSet.numGenes;
5 numLambdaVals = numel( lambdaVals );
6 cor = Ref.Coral;
7 ann = cor.Annotations;
8 identifierIndex = optionsSet.identifierIndex;
9 ids = ann.ids{ identifierIndex };
10 labels = ann.labels{ identifierIndex };
11 annot = get_annotation( cor, ann.identifier{ identifierIndex } );
12%focus on midbrain
13 anatInds = 10;
14 display( numGenes )
15 optionsSigned = struct( 'identifierIndex', identifierIndex, 'numGenesKept', numGenes, ...
16     'regionIndexInit', 2, 'lambdaMult', 0, 'saveResults', 0, ...
17     'verboseExec', 1, 'indsInAtlas', anatInds, 'positiveConstraint', 0 );
18 optionsPos = struct( 'identifierIndex', identifierIndex, 'numGenesKept', numGenes, ...
19     'regionIndexInit', 2, 'lambdaMult', 0, 'saveResults', 0, ...
20     'verboseExec', 1, 'indsInAtlas', anatInds, 'positiveConstraint', 1 );
21 for vv = 1 : numLambdaVals
22     lambdaCurrent = lambdaVals( vv )
23     if lambdaCurrent > 0
24         optionsSigned.lambdaMult = lambdaCurrent;
25         optionsPos.lambdaMult = lambdaCurrent;
26         signedFitting = signed_fitting( Ref, D, optionsSigned );
27         positiveFitting = signed_fitting( Ref, D, optionsPos );
28         exploreL1L2ParameterSpaceLoc = struct( 'signedFitting', signedFitting, ...
29             'positiveFitting', positiveFitting );
30         exploreL1L2ParameterSpace{ vv } = exploreL1L2ParameterSpaceLoc;
31         for aa = 1 : numel( anatInds )
32             anatInd = anatInds( aa );
33             signedFittingCoeffs = signedFitting.coeffsReg{ anatInd };
34             numGenesKept = signedFitting.numGenesKept;
35             signedFittingIndsKept = signedFitting.fittingInds{ anatInd };
36             signedFittingIndsKept = signedFittingIndsKept( 1 : numGenesKept );
37             signedFittingSol = D( :, signedFittingIndsKept ) * ...
38                 signedFittingCoeffs;
39             signedFittingSol = make_volume_from_labels( signedFittingSol, ...
40                         brainFilter );

```

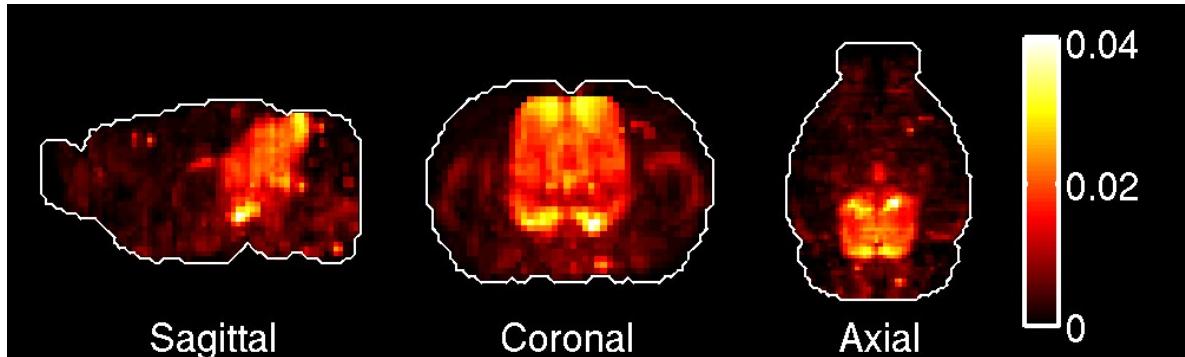


Figure 2.6: Maximal-intensity projection of the best-fitted set of genes (with weights of both signs) in the midbrain (constructed out of the best-fitted 100 genes)

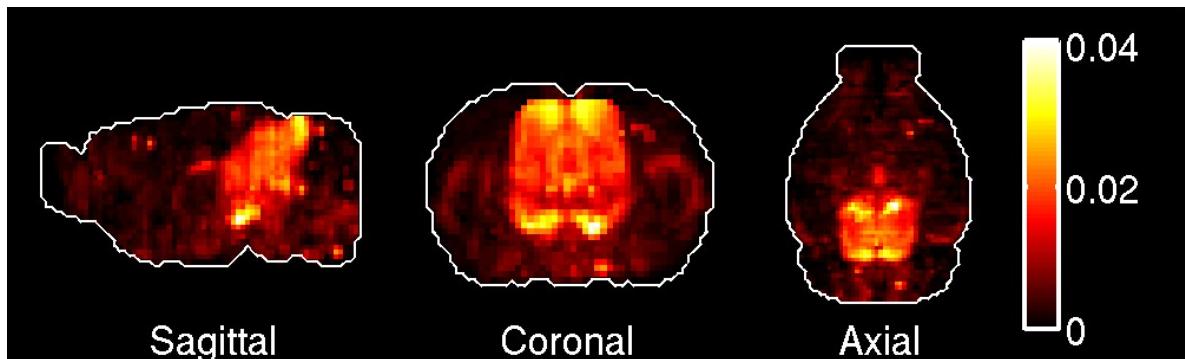


Figure 2.7: Maximal-intensity projection of the best-fitted set of genes (with positive weights) in the midbrain (constructed out of the best-fitted 100 genes).

```

41 plot_intensity_projections( signedFittingSol );
42 pause( 2 );
43 positiveFittingCoeffs = positiveFitting.coeffsReg{ anatInd };
44 numGenesKept = signedFitting.numGenesKept;
45 positiveFittingIndsKept =...
46     positiveFitting.fittingInds{ anatInd }( 1 : numGenesKept );
47 positiveFittingSol = D( :, positiveFittingIndsKept ) *...
48     signedFittingCoeffs;
49 positiveFittingSol = make_volume_from_labels( positiveFittingSol ,...
50     brainFilter );
51 plot_intensity_projections( positiveFittingSol );
52 pause( 2 );
53 end
54 end
55 end

```

Chapter 3

Genes versus genes: co-expression networks

3.1 Co-expression network of genes in the Allen Atlas

The co-expression of two genes in the Allen Atlas is defined as the cosine similarity between their gene-expression vectors in voxel space. Given a voxel-by-gene matrix containing the brain-wide expression energies (as in Equation 1.1), the corresponding gene-by-gene matrix of co-expressions of the full set of genes, or $\text{coExpr}^{\text{full}}$, is the symmetric matrix with entries equal to the co-expression of pairs of genes, as in Equation 3.1.

$$\text{coExpr}^{\text{full}}(g, g') := \frac{\sum_{v=1}^V E(v, g)E(v, g')}{\sqrt{\sum_{u=1}^V E(u, g)^2 \sum_{w=1}^V E(w, g')^2}}. \quad (3.1)$$

The matrix $\text{coExpr}^{\text{full}}$ can be computed as follows in Matlab:

```
1cor = Ref.Coronal;
2% divide each column of the data matrix by its L2-norm
3DNormalised = normalise_integral_L2(D);
4% co-expression matrix of the full set of genes
5coExpressionFull = DNormalised' * DNormalised;
```

The function `co_expression_matrix.m` takes a voxel-by-gene matrix as an argument and returns the gene-by-gene co-expression matrix defined by Equation 3.1, which equals the matrix `coExpressionFull` defined in the above code snippet if the full voxel-by-gene matrix of gene expression energies is used as an argument. Other versions of the Allen Atlas than the brain-wide standard annotations can be specified in the options to restrict the voxels to one of the annotations described in Table 1.1.

- **Example 18. Distribution of co-expression coefficients.** The diagonal elements of the co-expression matrix equal 1 by construction, and the co-expression matrix is symmetric. Hence, the distribution of co-expression coefficients in the atlas is given by the upper diagonal coefficients of the co-expression matrix, which can be extracted using the function `upper_diagonal_coeffs.m`, as in the code snipped below, which plots the distribution of brain-wide co-expression coefficients (Figure 3.1)) and compares it to the one of co-expression coefficients in the left hemisphere (as defined in the `big12` annotation, Figure 3.2).

```

1coExpressionFull = co_expression_matrix( Ref, D );
2% restrict the gene-expression data to the voxels that are
3% in the big12 annotation
4 optionsLeft = struct( 'identifierIndex', 5 );
5% cosine distances between the gene-expression vectors of genes
6% in the left hemisphere (voxels annotated in the big12 version of the ARA)
7coExpressionLeft= co_expression_matrix( Ref, D, optionsLeft );
8% consider only the non-trivial co-expressions
9upperCoeffsFull = upper_diagonal_coeffs( coExpressionFull );
10upperCoeffsLeft = upper_diagonal_coeffs( coExpressionLeft );
11% sort the brain-wide co-expression coefficients
12[ valsFull, inds ] = sort( upperCoeffsFull );
13% plot the brain-wide co-expression coefficients
14coFigure = figure( 'Color', 'w', 'InvertHardCopy', 'off',...
15 'Position', [ 200, 200, 800, 600 ] );
16plot( valsFull, '.b', 'markersize', 4 );
17xlabel( 'Sorted pairs of genes', 'fontsize', 20, 'fontweight', 'b' );
18ylabel( 'Brain-wide co-expression', 'fontsize', 20, 'fontweight', 'b' );
19set( gca, 'YLim', [ 0, 1 ], 'dataAspectRatio', [ 3*10^6 1 1 ],...
20 'fontsize', 20, 'fontweight', 'b' );
21 pause;
22hold off;
23close all;
24% plot upperCoeffsLeft against upperCoeffsFull
25figure( 'Color', 'w', 'InvertHardCopy', 'off',...
26 'Position', [ 200, 200, 800, 800 ] );
27plot( valsFull, upperCoeffsLeft( inds ), '.r', 'markersize', 1 );
28xlabel( 'Co-expression in the left hemisphere',...
29 'fontsize', 20, 'fontweight', 'b' );
30ylabel( 'Brain-wide co-expression', 'fontsize', 20,...
31 'fontweight', 'b' );
32set( gca, 'YLim', [ 0, 1 ], 'dataAspectRatio', [ 1 1 1 ],...
33 'fontsize', 20, 'fontweight', 'b' );

```

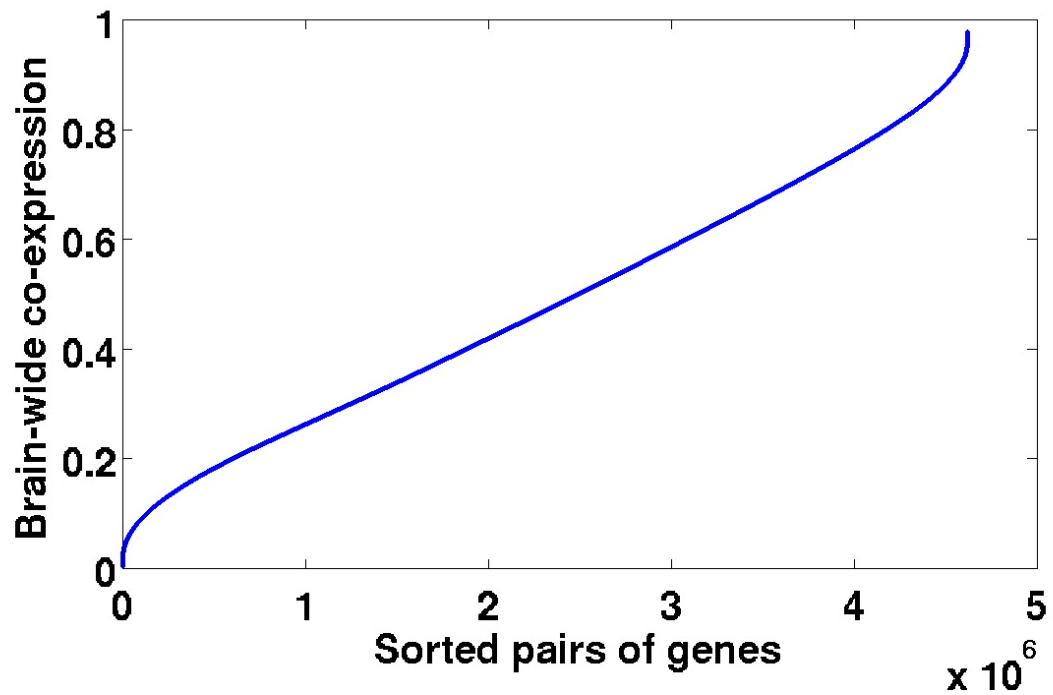


Figure 3.1: Sorted upper-diagonal elements of the brain-wide co-expression matrix.

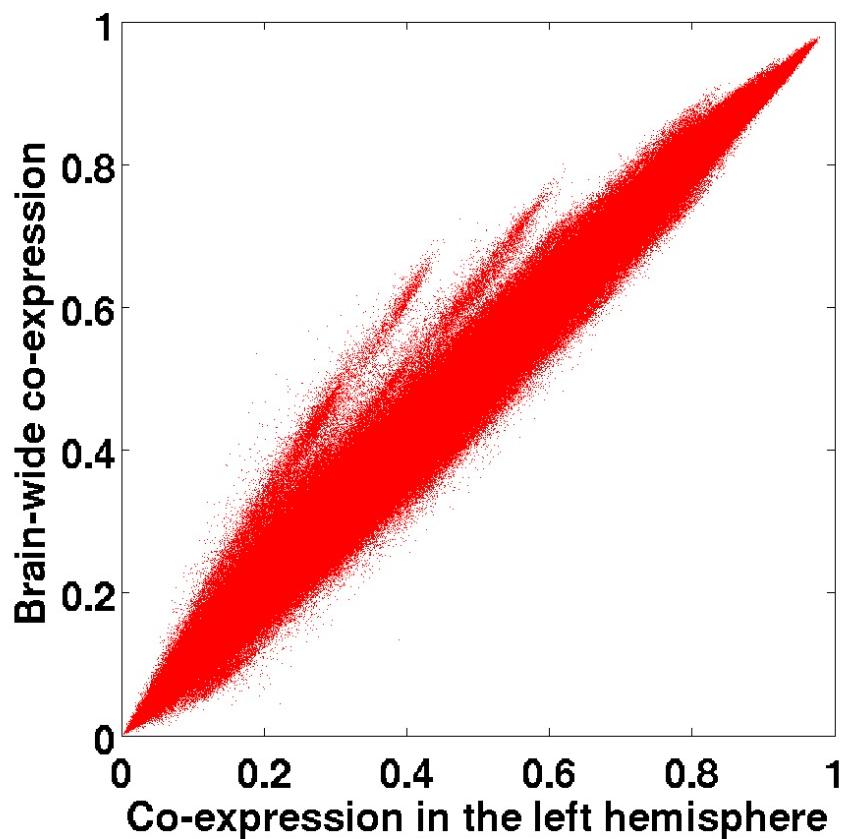


Figure 3.2: Upper-diagonal elements of the co-expression matrix in the left hemisphere plotted against upper-diagonal elements of the brainwide co-expression matrix. The deviation from the diagonal reflects the difference between the sets of voxels in the `big12` and `standard` annotations

3.2 Monte Carlo analysis of brain-wide co-expression networks

3.2.1 Special sets of genes versus full atlas

The gene-by-gene matrix $\text{coExpr}^{\text{full}}$ defines a universe in which we would like to study co-expression networks of special sets of genes, in a probabilistic way.

$$\text{coExpr}^{\text{full}} \in \mathcal{M}_{G_{\text{full}}}(\mathbf{R}), \quad (3.2)$$

Given a set of G_{special} genes of interest, corresponding to the column indices $(g_1, \dots, g_{\text{special}})$ in the data matrix, their co-expression matrix $\text{coExpr}^{\text{special}}$ is obtained by extracting the submatrix of $\text{coExpr}^{\text{full}}$ corresponding to these indices (see Equation 3.2.4).

$$\text{coExpr}^{\text{special}} \in \mathcal{M}_{G_{\text{special}}}(\mathbf{R}), \quad (3.3)$$

$$\text{coExpr}^{\text{special}}(i, j) = \text{coExpr}^{\text{full}}(g_i, g_j), \quad i, j \in [1..G_{\text{special}}]. \quad (3.4)$$

3.2.2 Cumulative distribution function of co-expression coefficients in sets of genes drawn from the Allen Brain Atlas

Having observed (Figure 3.1) that the distribution of pairwise co-expression coefficients of genes in the whole coronal atlas is roughly linear in a large domain of co-expression, we can study the cumulative distribution function of the co-expression coefficients in the special set, and compare it to the one resulting from random sets of genes (with the same number of genes as the special set, in order to eliminate the sample-size bias).

These cumulative distribution functions are evaluated in the following way. Again let G_{special} denote the size of the matrix $\text{coExpr}^{\text{special}}$, i.e. the number of genes from which $\text{coExpr}^{\text{special}}$ was computed. Consider the set of entries above the diagonal of $\text{coExpr}^{\text{special}}$ above the diagonal (which are the meaningful quantities in $\text{coExpr}^{\text{special}}$):

$$C^{\text{special}} = \{\text{coExpr}^{\text{special}}(g, h), 1 \leq g \leq G_{\text{special}}, h > g\}. \quad (3.5)$$

The elements of this set are numbers between 0 and 1. For every number between 0 and 1, the cumulative distribution function (c.d.f.) of C^{special} , denoted by cdf^C is defined as the fraction of the elements of C^{special} that are smaller than this number:

$$\begin{aligned} \text{cdf}^{\text{special}} : [0, 1] &\rightarrow [0, 1] \\ x &\mapsto \frac{1}{|C^{\text{special}}|} \sum_{c \in C^{\text{special}}} \mathbf{1}(c \leq x), \end{aligned} \quad (3.6)$$

where $|C^{\text{special}}| = G_{\text{special}}(G_{\text{special}} - 1)/2$.

For any set of genes, $\text{cdf}^{\text{special}}$ is a growing function $\text{cdf}^{\text{special}}(0) = 0$ and $\text{cdf}^{\text{special}}(1) = 1$. For highly co-expressed genes, the growth of $\text{cdf}^{\text{special}}$ is concentrated at high values of the argument (in the limit where all the genes in the special set have the same brain-wide expression vector, all the entries of the co-expression matrix go to 1 and the cumulative distribution function converges to a Dirac measure supported at 1). To compare the function $\text{cdf}^{\text{special}}$ to what could be expected by chance, let us draw R random sets of G_{special} genes from the Atlas, compute their co-expression network by extracting the corresponding entries from the full co-expression matrix of the atlas ($\text{coExpr}^{\text{full}}$). This induces a family of R growing functions $\text{cdf}_i, 1 \leq i \leq R$ on the interval $[0, 1]$

$$\begin{aligned} \text{cdf}_i &: [0, 1] \rightarrow [0, 1], 1 \leq i \leq R \\ \text{cdf}_i(0) &= 0, \quad \text{cdf}_i(1) = 1. \end{aligned} \tag{3.7}$$

From this family of functions, we can estimate a mean cumulative distribution function $\langle \text{cdf} \rangle$ of the co-expression of sets of G_{special} genes drawn from the Allen Atlas, by taking the mean of the values of cdf_i across the random draws:

$$\forall x \in [0, 1], \quad \langle \text{cdf} \rangle(x) = \frac{1}{R} \sum_{i=1}^R \text{cdf}_i(x). \tag{3.8}$$

Standard deviations cdf^{dev} of the distribution of c.d.f.s are estimated in the same way on the interval $[0, 1]$ (which the user of the code can discretize into a regular grid using the `coExprStep` component of the options of the function `cumul_co_expr.m`, see example below):

$$\forall x \in [0, 1], \quad \text{cdf}^{\text{dev}}(x) = \sqrt{\frac{1}{R} \sum_{i=1}^R (\text{cdf}_i(x) - \langle \text{cdf} \rangle(x))^2}. \tag{3.9}$$

The functions $\langle \text{cdf}^{\text{special}} \rangle$, $\langle \text{cdf} \rangle$ and cdf^{dev} are fields of the output of the function `cumul_co_expr.m` whose usage is illustrated in the example below.

Example 19. Cumulative distribution function of co-expression of a special set of genes. Consider the set of 288 genes from the NicSNP database, whose position in the data matrix is encoded in the Matlab file `nicotineGenes.mat`. NB: the code snippet below can be applied to any special set of genes upon changing the variable `indsSmall`. It should reproduce Figure 3.3.

```

1 load('nicotineGenes.mat');
2 indsSmall = nicotineGenes.nicotineIndicesInTop75;
3 % extract the special co-expression network
4 coExpressionSmall = coExpressionFull(indsSmall, indsSmall);
5 % compute the c.d.f., and simulate the distribution of

```

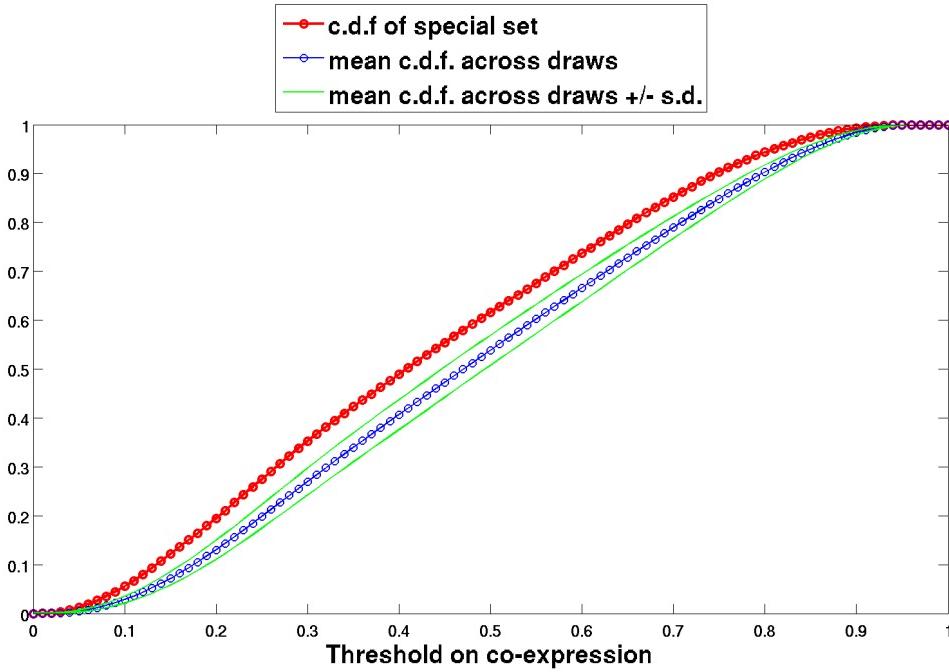


Figure 3.3: Cumulated distribution functions of the upper-diagonal entries of the co-expression matrix of the special set of genes listed in `nicotineGenes.mat`.

```

6% c.d.f.s across numPaths random sets of genes of size numel( indsSmall )
7numPaths = 1000;
8optionsCDF = struct( 'numPaths', numPaths, 'coExprStep', 0.01, ...
9      'minCoExpr', 0, 'maxCoExpr', 1, 'wantPlots', 1, 'species', 'mouse' );
10cumulCoExpr = cumul_co_expr( coExpressionFull, coExpressionSmall, optionsCDF );
11% plot the results
12optionsCDF = struct( 'savePlots', 1, 'fontSize', 18, ...
13  'figurePosition', [200 200 1100 700], 'markerSize', 14, ...
14      'lineWidth', 2, 'fileName', 'cumulCoExprFunction' );
15cumulCoExprPlot = cumul_co_expr_plot( cumulCoExpr, optionsCDF );

```

3.2.3 Thresholding the co-expression matrix

The co-expression matrix $\text{coExpr}^{\text{special}}$ corresponding to a special subset of the genes in the Allen Atlas (Equation 3.2.4) is symmetric, like $\text{coExpr}^{\text{full}}$, and its entries are in the interval $[0, 1]$. It can be mapped to a weighted graph in the following way (see [5] for details). The vertices of the graph are the genes, and the edges are as follows:

- genes g and g' are linked by an edge if their co-expression is strictly positive.
- If an edge exists, it has weight $\text{coExpr}_{gg'}^{\text{special}}$.

Let us define the following thresholding procedure on co-expression graphs: given a threshold ρ between 0 and 1, put to zero all the entries of $\text{coExpr}^{\text{special}}$ that are lower than this coefficient. The underlying graph is obtained by taking the graph corresponding to $\text{coExpr}^{\text{special}}$, and cutting all the links with weight below ρ .

$$\text{coExpr}_{\rho}^{\text{special}}(g, h) = \text{coExpr}^{\text{special}}(g, h) \times \mathbf{1}(\text{coExpr}^{\text{special}}(g, h) \geq \rho). \quad (3.10)$$

The more-co-expressed a set of genes is, the larger the connected components of the thresholded graphs underlying $\text{coExpr}_{\rho}^{\text{special}}$ will be, for any value of the threshold ρ . For instance we can study the average size of connected components of thresholded co-expression matrices and the size of the largest connected component as a function of the threshold ρ :

$$\mathcal{A}(\rho) = \frac{\sum_{k=1}^G k N_{\rho}(k)}{\sum_{k=1}^G N_{\rho}(k)}, \quad (3.11)$$

$$\mathcal{M}(\rho) = \max \{k \in [1..G], N_{\rho}(k) > 0\}, \quad (3.12)$$

where $N_{\rho}(k)$ is the number of connected components with size k .

3.2.4 Statistics of sizes of connected components

For any quantity worked out from the special co-expression matrix co-expression matrix $\text{coExpr}^{\text{special}}$ defined in Equation , we can simulate its probability distribution by repeatedly drawing random random sets of genes from the atlas, and recomputing the same quantity on for this set.

Let us use the above-defined thresholding procedure to study a set of $G_{\text{special}} = 288$ genes obtained by intersecting the NicSNP database [6] with the set of $G = 3,041$ genes given by

`get_genes(Ref.Coral, 'top75CorrNoDup', 'allen').` We would like to ascertain whether this set of genes is more co-expressed than expected by chance for a set of this size taken from `genesAllen`. At each value of a regular grid the threshold ρ between zero and 1, the function `co_expression_island_bootstrap.m` computes the maximal size and average size of connected components of the thresholded co-expression graph, and draws R random sets of genes of size G_{special} from the atlas. This induces a distribution of R

Example 20. Consider the set of 288 genes from the NicSNP database, whose position in the data matrix is encoded in the Matlab file `nicotineGenes.mat`.

```
% compute the full co-expression matrix coExpressionFull
2coExpressionFull = co_expression_matrix( D );
3load( 'nicotineGenes.mat' );
4indsSmall = nicotineGenes.nicotineIndicesInTop75;
5coExpressionSmall = coExpressionFull( indsSmall, indsSmall );
% Monte Carlo analysis of the graph underlying the co-expression matrix
```

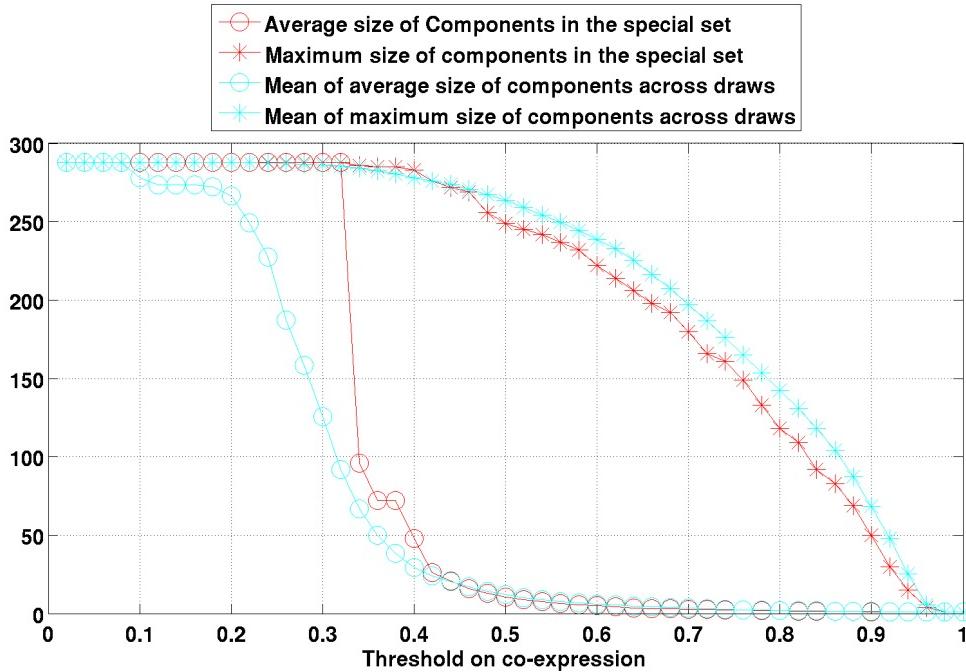


Figure 3.4: Monte Carlo analysis of the graph underlying the co-expression matrix of 288 genes from the NicSNP database. Average and maximum size of connected components as a function of the threshold.

```

7 numDraws = 1000;
8 optionsBootstrap = struct( 'thresholdInit', 1, 'thresholdStep', thresholdStep, ...
9   'numDraws', numDraws, 'thresholdFinal', min( min( coExpressionFull ) ) );
10 coExpressionComponentsBootStrap =...
11   co_expression_components_bootstrap( coExpressionFull, ...
12     coExpressionSmall, optionsBootstrap );
13 % reproduce the three plots of this example and save them
14 optionsPlot = struct( 'savePlots', 1, ...
15   'fontSize', 16, 'figurePosition', [200 200 1100 700], 'markerSize', 14 );
16 fileNames = { 'coExpressionSet', 'coExpressionProba', 'coExpressionSDs' };
17 coExpressionComponentsBootStrapPlot =...
18   co_expression_components_bootstrap_plot(
19 coExpressionComponentsBootStrap, optionsPlot, fileNames );

```

3.2.5 Neuroanatomical properties of connected components

Relation between fitting scores and co-expression

If the co-expression of pairs of genes is defined as the cosine of the angle between their expression vectors in voxel space, as in Equation 3.1, the fitting score of a gene g to a

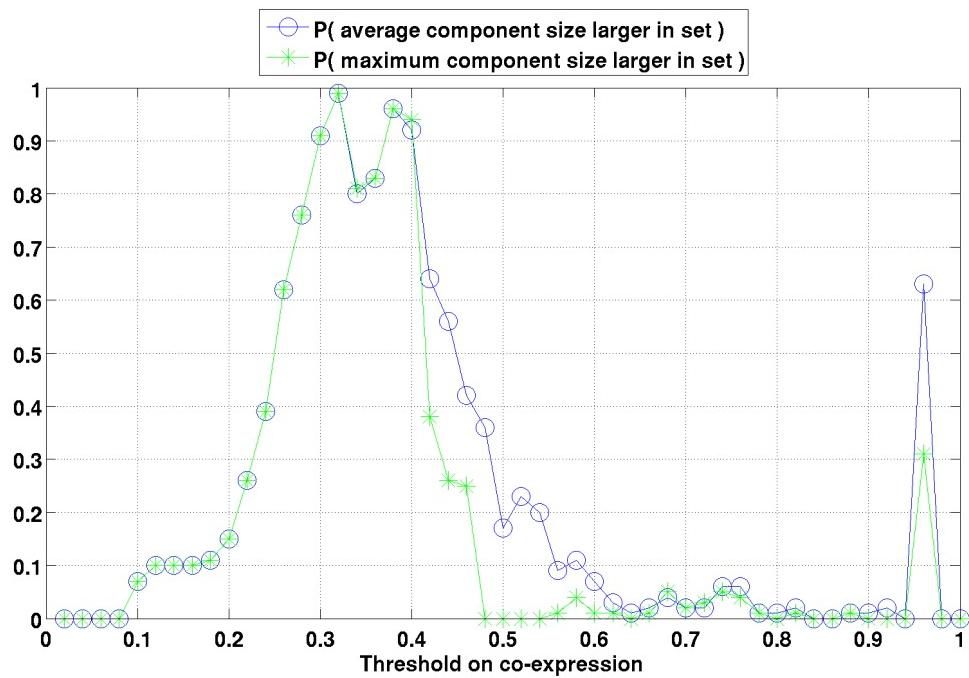


Figure 3.5: Monte Carlo analysis of the graph underlying the co-expression matrix of 288 genes from the NicSNP database. Estimated probabilities for the average and maximum size of connected components to be larger than in random sets of genes of the same size.

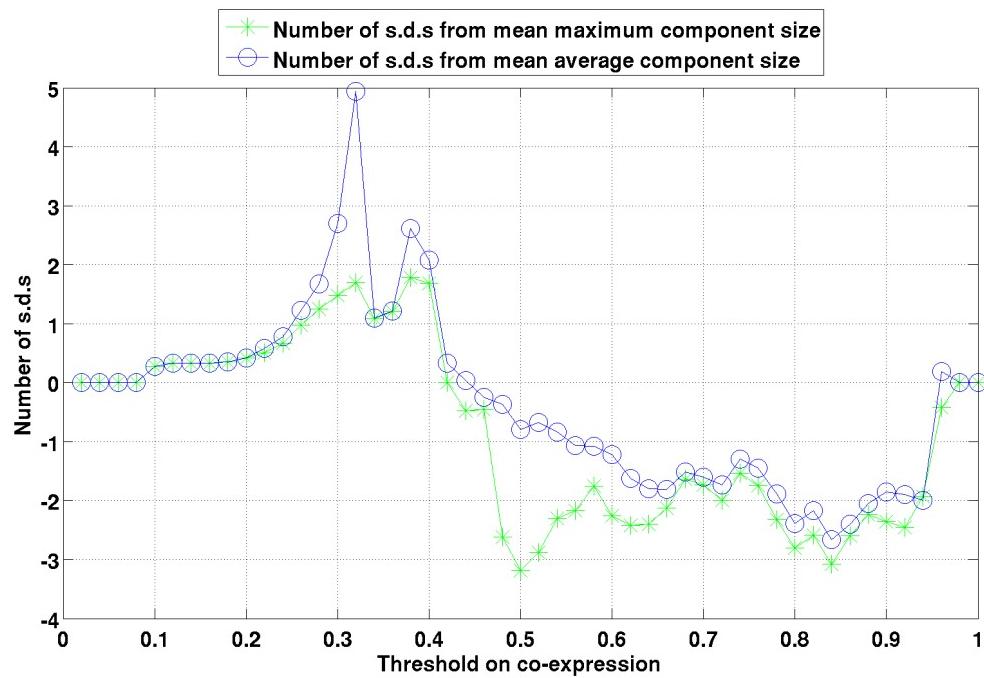


Figure 3.6: Monte Carlo analysis of the graph underlying the co-expression matrix of 288 genes from the NicSNP database. Deviation of average and maximum size of connected components from mean (expressed in number of standard deviations at every value of the threshold).

region ω of the brain equals the co-expression of gene g and a (hypothetical) gene whose expression profile would be proportional to the characteristic function of region ω :

$$\phi_\omega(g) = 1 - \frac{1}{2} \sum_{v \in \Omega} (E_g^{\text{norm}}(v) - \chi_\omega(v)) = \sum_{v \in \Omega} E_g^{\text{norm}}(v) \chi_\omega(v), \quad (3.13)$$

where the second equality comes from the normalization of E_g^{norm} and χ_ω in voxel space.

Fitting scores of sums of gene-expression vectors

At a given level of the threshold on co-expression, a set of genes is partitioned into connected components induced by the graph underlying the thresholed matrix defined in Equation . For a set of genes of fixed size `numGenes`, the function `fitting_distribution_in_atlas.m` estimates the distribution of fitting scores of the sum of sets of genes of a given size, extracted from the Allen Atlas.

In these functions, the genes are not weighted by coefficients to be optimized, they are simply summed over, so that the set of gene-expression vectors is different from the one explored in marker genes. The fitting score ϕ_ω^{sum} of the sum of genes depends only on a list of K distinct K genes $\{g_1, \dots, g_K\}$, and a region ω in the Allen Reference Atlas:

$$\phi_\omega(\{g_1, \dots, g_K\}) := \sum_{v \in \Omega} E_{\{g_1, \dots, g_K\}}^{\text{norm}}(v) \chi_\omega(v), \quad (3.14)$$

where $E^{\text{norm}}(\{g_1, \dots, g_K\})$ is the sum of gene-expression vectors of the genes $\{g_1, \dots, g_K\}$, normalized in the L^2 sense:

$$E_{\{g_1, \dots, g_K\}}^{\text{norm}}(v) = \frac{\sum_{i=1}^K E(v, g_i)}{\sqrt{\sum_{v \in \Omega} \left(\sum_{j=1}^K E(v, g_j) \right)^2}}. \quad (3.15)$$

Example 21. Fitting score of a sum of genes in a given region. It can be computed using the same functions as for the original data matrix. The sum of the relevant columns of the data matrix has to be substituted to the second argument of the function `fitting_from_id.m`:

```

1 colIndsForGenes = 1 : 10;
2 % the sum of these gene expressions (as a column vector )
3 sumOfDataForGenes = ( sum( ( D( :, colIndsForGenes ) )' ) );
4 % consider the big12 annotation
5 identifierIndex = 5;
6 fittingScoresInAtlasBig12 = fitting_region_by_gene( Ref, sumOfDataForGenes, ...
7 identifierIndex);

```

Example 22. Neuroanatomy of nicotine-related genes in the big12 annotation.

```

1load( 'nicotineGenes.mat' );
2indsSmall = nicotineGenes.nicotineIndicesInTop75;
3coExpressionSmall = coExpressionFull( indsSmall, indsSmall );
4
5% compute the connected components of the thresholded network
6optionsComponents = struct( 'thresholdInit', 1, 'thresholdFinal', 0, ...
7  'thresholdStep', 0.02 );
8coExpressionComponents = co_expression_components( coExpressionSmall, ...
9  optionsComponents );
10% compare the fitting scores of the connected components
11% to the estimated distribution
12% of fitting scores of sets of the same size
13optionsAnatomy = struct( 'minimalComponentSize', 2, 'identifierIndex', 5, ...
14  'numDraws', 100 );
15coExprComponentsAnatomy = co_expr_components_anatomy( Ref, D, ...
16  coExpressionComponents, indsSmall, optionsAnatomy );

```

One can search the results by P -value of fitting scores, and/or size of connected components, and/or brain region:

Example 23. Neuroanatomy of nicotine-related genes in the big12 annotation (continued).

```

1% show all the connected components, at any value of
2% the threshold, whose fitting score to any region is estimated to be in
3%
4probaCrit = 0.001;
5extractComponentByPValue = extract_component_by_p_value( Ref, D, ...
6  coExprComponentsAnatomy, probaCrit );
7% illustrate the sums of these components
8% and their anatomical properties
9numComponentsExtracted = numel( extractByPvalue.threshold );
10display( numComponentsExtracted )
11if numComponentsExtracted > 0
12  for kk = 1 : numComponentsExtracted
13    figureComp = figure_for_component( Ref, D, extractByPvalue, kk );
14    pause;
15  end
16end

```

Chapter 4

Genes versus cell types (to be released in the second version of the toolbox, not publicly available as of November 2012)

4.1 Cell-type-specific microarray data

The file `G_t_means.txt` contains a matrix of microarray data. The rows correspond to the cell-type-specific samples coming from several studies and analyzed in [7]. The columns correspond to genes (the matrix is a type-by-gene matrix). The following code snippet (included in the file `cell_type_start_up.m`) defines two matrices with the same numbers of columns, corresponding to the intersection of the sets of genes in the coronal Allen Brain Atlas and in the microarray data:

```
% load the cell-type-specific data
2C = dlmread( 'G_t_means.txt' );
% which columns of the data matrix C
% correspond to genes that are also in D?
5load( 'colsToUseInAllen.mat' );
6load( 'colsToUseInTypes.mat' );
%construct the corresponding matrices
8DUsed = D( :, colsToUseInAllen );
9CUsed = C( :, colsToUseInTypes );
10% Names of the cell types, in the same order as the rows of the matrix CUsed
11load( 'cellTypesDescription.mat' );
```

4.2 Brain-wide correlation profiles

The function `cell_types_correls.m` computes the voxel-by-type correlation matrix between the coronal Allen Brain Atlas and a set of cell types, as in equation 4.1.

$$\text{Corr}(v, t) = \frac{\sum_{g=1}^G (C(t, g) - \bar{C}(g))(E(v, g) - \bar{E}(g))}{\sqrt{\sum_{g=1}^G (C(t, g) - \bar{C}(g))^2} \sqrt{\sum_{g=1}^G (E(v, g) - \bar{E}(g))^2}}, \quad (4.1)$$

$$\bar{C}(g) = \frac{1}{T} \sum_{t=1}^T C(t, g), \quad (4.2)$$

$$\bar{E}(g) = \frac{1}{V} \sum_{v=1}^V E(v, g). \quad (4.3)$$

Example 24. Brain-wide correlation profiles between cell-type-specific data and the Allen Atlas. To compute the correlations for all the available cell types, the type indices specified in the variables span the whole range of row indices of the matrix `C`:

```
% C is a type-by-gene matrix
2numTypes = size( C, 1 );
3typeIndices = 1 : numTypes;
4cellTypesCorrelations = cell_types_correls( Ref, D, C, ...
5    colsToUseInAllen, colsToUseInTypes, typeIndices );
```

The results are saved in the file `cellTypesCorrelations.mat`. It has $V = 49,742$ rows and $T = 64$ columns (the matrix `cellTypesCorrelations` is a voxel-by-type matrix, consistently with the definition of $\text{Corr}(v, t)$ in Equation 4.1), so that each column can be mapped to a volume and visualized in exactly the same way as a column of the data matrix `D`:

Example 25. Visualization of brain-wide correlation profiles. Let us plot maximal-intensity projections of the correlations between the 20-th row of C and all the rows of the Allen Atlas.

```
% load the correlations, unless of course they have just been
% computed using the previous example
3load( 'cellTypesCorrelations.mat' );
% take a column of the correlation matrix,
% i.e. focus on a specific cell type
6typeIndex = 20;
% describe the cell type specified by typeIndex
8display( cellTypesDescription( typeIndex ) )
9correlsTypeToAtlas = cellTypesCorrelations( :, typeIndex );
```

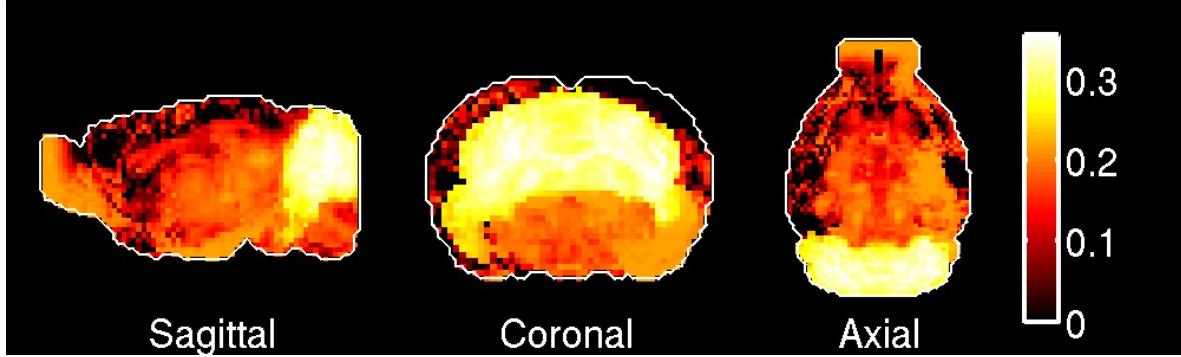


Figure 4.1: Maximal-intensity projections of correlations between the voxels of the coronal Allen Atlas ($G=3,041$ genes) and granule cells.

```

10% map the column to a three-dimensional grid
11correlsTypeToAtlasVol = make_volume_from_labels( correlsTypeToAtlas , ...
12                                brainFilter );
13plot_intensity_projections( correlsTypeToAtlasVol );

```

4.3 Brain-wide density estimates of cell types

To decompose the gene-expression at every voxel in the Allen Atlas into its cell-type-specific components, let us introduce the positive quantity $\rho_t(v)$ denoting the contribution of cell-type t at voxel v , and propose the following linear model:

$$E(v, g) = \sum_{t=1}^T \rho_t(v) C_t(g) + \text{Residual}(v, g). \quad (4.4)$$

Both sides are estimators of the number of mRNAs for gene g at voxel v . The residual term in Equation 4.4 reflects the fact that $T = 64$ cell types are not enough to sample the whole diversity of cell types in the mouse brain, as well as noise in the measurements, reproducibility issues, the non-linearity of the relations between numbers of mRNAs, expression energies and microarray data.

To find the best fit of the model, we have to minimize the residual term by solving the following problem :

$$(\rho_t(v))_{1 \leq t \leq T, 1 \leq v \leq V} = \operatorname{argmin}_{\phi \in \mathbf{R}_+(T, V)} \mathcal{E}_{E, C}(\phi), \quad (4.5)$$

where

$$\mathcal{E}_{E, C}(\phi) = \sum_{v=1}^V \left(\sum_{g=1}^G \left(E(v, g) - \sum_{t=1}^T \phi(t, v) C_t(g) \right)^2 \right). \quad (4.6)$$

As the terms of the sum over voxels that involve a fixed voxel v involve only terms of the test function ϕ with the same voxel index, and are positive, the problem 4.5 can be solved voxel by voxel. For each voxel we have to minimize a quadratic function of a vector with T positive components.

$$\forall v \in [1..V], (\rho_t(v))_{1 \leq t \leq T} = \operatorname{argmin}_{\nu \in \mathbf{R}_+^T} \sum_{g=1}^G \left(E(v, g) - \sum_{t=1}^T \nu(t) C_t(g) \right)^2. \quad (4.7)$$

For each cell type t , the coefficients $(\rho_t(v))_{1 \leq v \leq V}$ yield a brain-wide fitting profile between this cell type and the Allen Atlas.

The function `fit_voxels_to_types_picked.m` solves the quadratic optimization problems of Equation 4.7, one per voxel, for a list of cell-type indices, and returns the result in a voxel-by-type format.

Example 26. Consider the full cell-type-specific dataset with `numTypes` cell types, and estimate the brain-wide density of each of these cell types:

```
1 numTypes = size(C, 1);
2 typeIndicesPicked = 1 : numTypes;
3 fitVoxelsToTypesFull = fit_voxels_to_types_picked(Ref, D, C, ...
4 colsToUseInAllen, colsToUseInTypes, typeIndicesPicked);
```

The result of the computation in the above code snippet are saved in the file `fitVoxelsToTypes.mat`. Note that the results depends on the set of types specified by `typeIndicesPicked`, since the corresponding columns of C play the role of a basis on which the gene-expression vector at each voxel is decomposed under positive constraints.

4.4 Cell-type-based computational neuroanatomy

As the matrix of cell-type-specific microarray data is presented in the type-by-gene format, a given cell type in this dataset corresponds to a row index t in $[1, T]$. The corresponding column of any voxel-by-type matrix (such as correlations `Corr`, Equation 4.1, or densities ρ , Equation 4.4) can be mapped to volumes using the function `make_volume_from_labels.m`. Projections of these volume can be visualized using the function `plot_intensity_projections.m`. One can supplement these projections with sections of these volumes. The sections can be chosen computationally as follows by ranking the brain regions in the Allen Reference Atlas using the brain-wide correlation or density profiles.

Ranking of regions by correlations with a cell type. Given a three-dimensional grid containing brain-wide correlations between a given cell type t and the Allen Atlas

(as defined in Equation 4.1), one can rank the regions in a given annotation of the Allen Reference Atlas by computing the average correlation across in each of the regions:

$$\overline{\text{Corr}}(r, t) = \frac{1}{|V_r|} \sum_{v \in V_r} \text{Corr}(v, t). \quad (4.8)$$

For each cell type t , each region r in the ARA is ranked according to its average correlation. Call its rank $\chi(r)$. In particular, the correlation profile yields a top region $\rho^\chi(t)$ in the ARA, for which $\chi(\rho^\chi(t)) = 1$. This is the region whose voxels are most highly correlated to cell type t on average.

Ranking of regions by estimated density of a cell type. Consider a system of annotation of the ARA, taken from Table 1.1 `Big12` is the default option in the Matlab implementation). For each cell-type index t and each region V_r (where r is an integer index taking values in $[1..R]$, where R is the number of regions in the chosen system of annotation), we can compute the contribution of the voxels of the region to the brain-wide density profile of cell type t :

$$\bar{\rho}(r, t) = \frac{1}{\sum_{w \in \text{Brain}} \rho_t(w)} \sum_{v \in V_r} \rho_t(v). \quad (4.9)$$

Given a cell type t , the regions in the ARA is ranked according to its average correlation. Call its rank $\phi(r)$. In particular, the correlation profile yields a top region $\rho^\phi(t)$ in the ARA, for which $\phi(\rho^\phi(t)) = 1$. This is the region whose voxels bring the largest total contribution to the vector ρ_t .

These two rankings are implemented in the function `classify_pattern.m` (see example below). The values of the fraction of total density and average correlation for a given cell type can be plotted (in the `big12` version of the ARA), using¹ the functions `plot_correlations_for_big12.m` and `plot_densities_for_big12.m`.

Example 27. Rank the regions of the `big12` annotations by average correlation and fraction of the density of cell types, and display the results for granule cells. The code snippet below should reproduce Figures 4.2 and 4.3.

```
% Rank brain regions (in the big12 annotation of the ARA)
% by their correlation and fitting profiles to each cell type
3 classifyOptions = struct('identifierIndex', 5);
4 classifyBig12 = classify_pattern(Ref, cellTypesCorrelations, ... ,
```

¹The labels used in these plots can be decoded by comparing the fields `Ref.Coronal.Annotations.labels5` and `Ref.Coronal.Annotations.symbols5`. This yields the following correspondences: Basic cell groups and regions = Brain, Cerebral cortex = CTX, Olfactory areas = OLF, Hippocampal region = HIP, Retrohippocampal region = RHP, Striatum = STR, Pallidum = PAL, Thalamus = TH, Hypothalamus = HY, Midbrain = MB, Pons = P, Medulla = MY, Cerebellum = CB.

```

5           fitVoxelsToTypes, classifyOptions );
6% the values of average correlation and fraction of density
7regionByTypeAvgCorrelation = classifyBig12.regionByTypeAverageCorrelations;
8regionByTypeFracDensity = classifyBig12.regionByTypeFractionFittings;
9% the ranks of the regions, in the same region-by-type format
10rankRegionsByCorrelation = classifyBig12.rankRegionsAverageCorrelations;
11rankRegionsByDensity = classifyBig12.rankRegionsFractionFittings;
12% Focus on the granule cells, index 20;
13typeIndex = 20;
14display( cellTypesDescription( typeIndex ) );
15granuleCellsAvgCorrelation = regionByTypeAvgCorrelation( :, typeIndex );
16granuleCellsCorrelationRanks = rankRegionsByCorrelation( :, typeIndex );
17granuleCellsFracDensity = regionByTypeFracDensity( :, typeIndex );
18granuleCellsRankRegionsByDensity = rankRegionsByDensity( :, typeIndex );
19identifierIndex = classifyOptions.identifierIndex;
20labels = ann.labels{ identifierIndex };
21granuleCellsRankedRegionsCorr = labels( granuleCellsCorrelationRanks );
22% regions in the big12 annotations, ranked by correlation
23display( granuleCellsRankedRegionsCorr );
24granuleCellsRankedRegionDensity = labels( granuleCellsRankRegionsByDensity );
25% regions in the big12 annotations, ranked by correlationdensity
26display( granuleCellsRankedRegionDensity );
27%confirm these rankings by plotting the underlying values
28plot_correlations_for_big12( Ref, granuleCellsAvgCorrelation );
29pause( 2 );
30hold off;
31plot_densities_for_big12( Ref, granuleCellsFracDensity );
32hold off;

```

Given a region in the ARA (our best guess according to one of the above criteria), we need to pick a section that intersects this region. The function `cell_type_vol_prepare.m` implements the choice of the section. It chooses the section (which can be sagittal, coronal or axial, specified by the field `options.sectionStyle` of the options) that intersect the desired region (specified by the fields `options.identifierIndex` and `options.regionIndexForSection`) along the largest number of voxels, unless the field `options.customIndex` equals 1. Then it works out the section (of the required style) at a position specified by `option.desiredIndex`. This function is used by the function `figure_for_types.m` to produce Figures like 4.4 and 4.5.

4.5 Visualization of correlation and density profiles of cell types

Example 28. Consider the pyramidal neurons, index 47 in the auxiliary dataset. The following code snippet computes the region in `big12` annotation of the ARA with the highest average correlation (as in Equation 4.8), and chooses the section of the brain through this section that maximizes the (section-wide) average correlation. It displays

4.5. VISUALIZATION OF CORRELATION AND DENSITY PROFILES OF CELL TYPES45

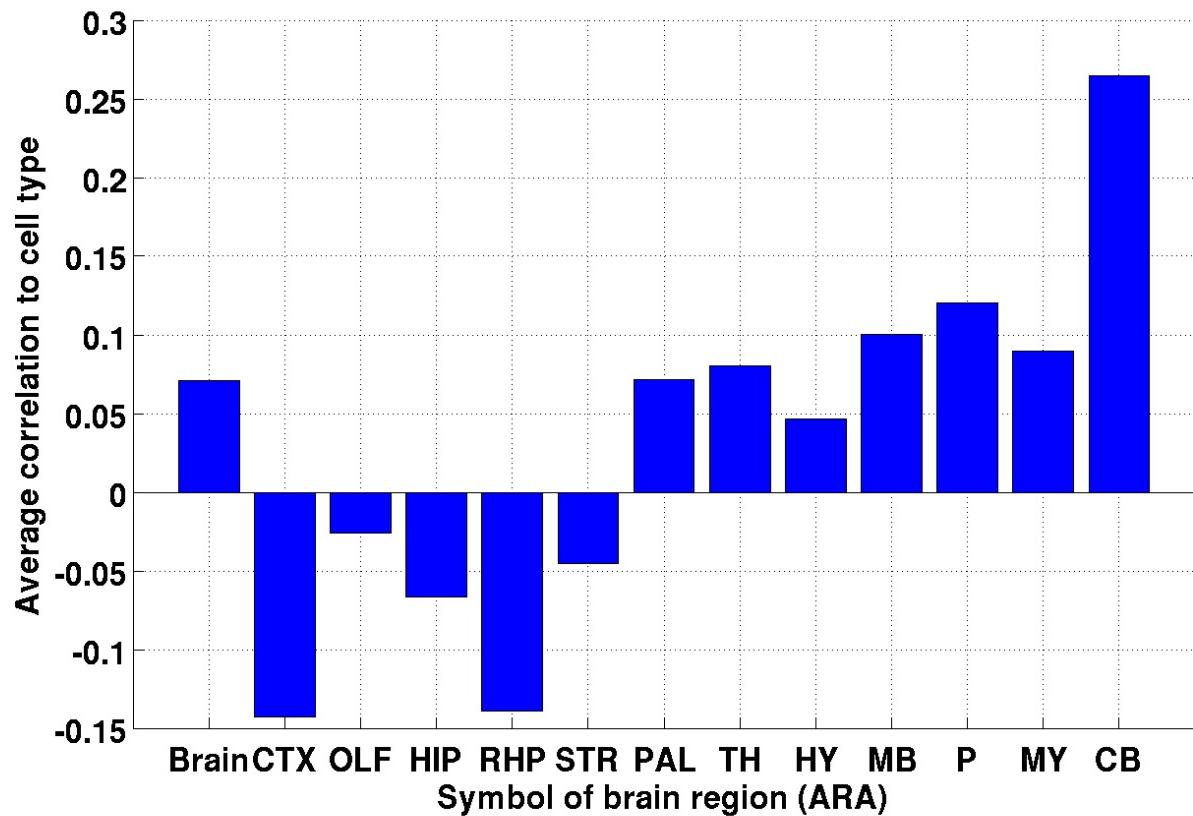


Figure 4.2: Average correlations between the voxels of the coronal Allen Atlas ($G=3,041$ genes) and granule cells (index 20 in the cell-type-specific dataset), in the regions of the big12 annotation of the ARA.

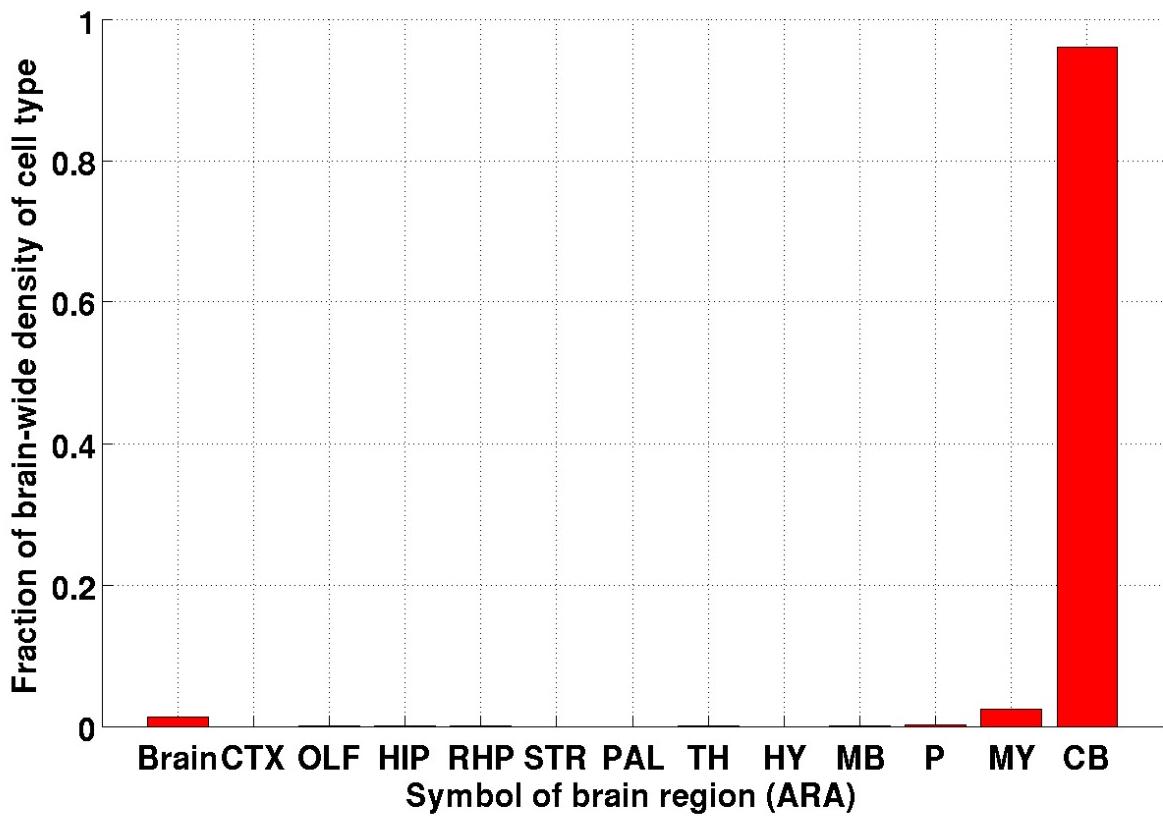


Figure 4.3: Fraction of the density of granule cells (index 20 in the cell-type-specific dataset) contributed by the regions in the `big12` annotation of the ARA.

4.5. VISUALIZATION OF CORRELATION AND DENSITY PROFILES OF CELL TYPES47

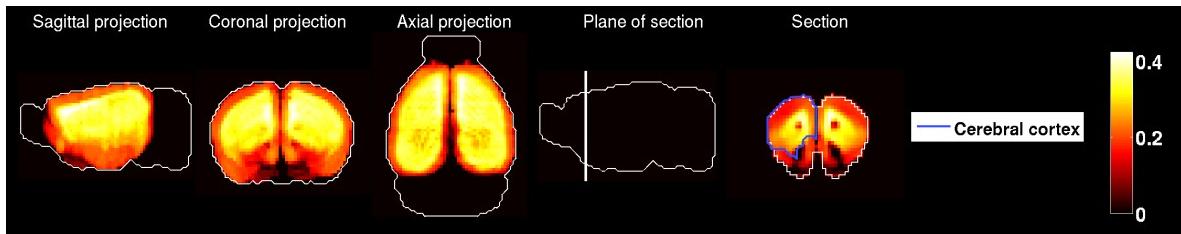


Figure 4.4: Correlations between the voxels of the coronal Allen Atlas ($G=3,041$ genes) and pyramidal cells (index 47 in the cell-type-specific dataset). Maximal-intensity projection and a section through the region in the ARA that maximizes the average correlation.

it next to the maximal-intensity projections of the correlation profiles, together with a graphical definition of the plane that was used to obtain the section. The snippet should reproduce Figure 4.4 and save it as `figureTypeCorrels47.png`.

```

1 typeIndex = 47;
2 display( cellTypesDescription( typeIndex ) )
3 correlsTypeToAtlas = cellTypesCorrelations( :, typeIndex );
4 correlsTypeToAtlasVol = make_volume_from_labels( correlsTypeToAtlas, ...
5 brainFilter );
6 optionsFig = struct( 'identifierIndex', 5, 'sectionStyle', 'coronal' );
7 figureForTypeCorrelations = figure_for_type_correlations( Ref, ...
8 correlsTypeToAtlasVol, optionsFig );

```

Example 29. Consider again the pyramidal neurons, index 47 in the auxiliary dataset. The following code snippet allows to visualize a section of their estimated brain-wide density, through the region of the ARA that contains the largest fraction of the total density, as per Equation 4.9. It should reproduce Figure 4.5 and save it as `figureTypeDensity47.png`.

```

1 typeIndex = 47;
2 densityOfType = fitVoxelsToTypes( :, typeIndex );
3 densityOfTypeVol = make_volume_from_labels( densityOfType, brainFilter );
4 figureForTypeDensity = figure_for_type_density( Ref, densityOfTypeVol, ...
5 optionsFig );
6 set( gcf, 'PaperPositionMode', 'auto' );
7 saveas( gcf, 'figureTypeDensity47.png', 'png' );

```

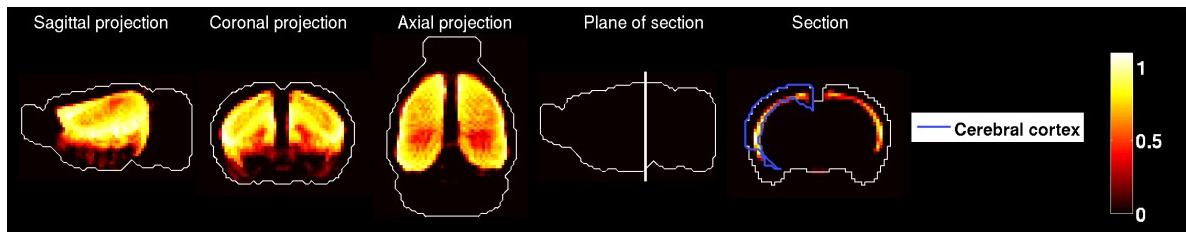


Figure 4.5: Brain-wide density estimate of pyramidal cells (index 47 in the cell-type-specific dataset). Maximal-intensity projection and a section through the region in the ARA that contains the largest fraction of total density.

Chapter 5

Clustering

5.1 Kullback–Leibler distance from gene-expression to a given probability distribution

One can use the Kullback–Leibler (KL) divergence to compare the brain-wide expression of a gene to a given probability distribution over the brain.

This probability distribution can be for instance the normalized average $E_{\text{full}}^{\text{avg}}$ across all genes in the coronal atlas for instance

$$\text{KL}^{\text{avg}}(g) := \sum_{v=1}^V \tilde{E}_g(v) \log \left(\frac{\tilde{E}_g(v)}{\tilde{E}_{\text{full}}^{\text{avg}}(v)} \right), \quad (5.1)$$

where $\tilde{E}_{\text{full}}^{\text{avg}}$ and \tilde{E}_g are probability densities over the brain obtained by normalization:

$$\tilde{E}_g(v) = \frac{E(v, g)}{\sum_{u=1}^V E(u, g)}, \quad \tilde{E}_{\text{full}}^{\text{avg}}(v) = \frac{E_{\text{full}}^{\text{avg}}(v)}{\sum_{u=1}^V E_{\text{full}}^{\text{avg}}(u)}. \quad (5.2)$$

5.2 Construction of a bipartite graph from a voxel-by-gene matrix

The Allen Atlas can be mapped to a *weighted bipartite graph* in the following way:

- the first set of vertices consists of voxels, numbered from 1 to V ,
- the second set of vertices consists of genes, numbered from 1 to G ,
- each of the edges connects one voxel, say v to one gene, say g and has a weight given by the expression energy of $E(v, g)$ of the gene at the voxel.

We looked for partitions of this weighted bipartite graph into subgraphs such that the weights of the internal edges of the subgraphs are strong compared to the weights of the edges between the subgraphs. This is the isoperimetric problem addressed by the algorithm of [8] (the graph need not be bipartite to apply this algorithm, but since we started with a bipartite graph, each of the subgraphs, or biclusters, returned by the algorithm, is bipartite, and therefore corresponds to a set of voxels and a set of genes).

Given a weighted graph, the algorithm cuts some of the links, thus partitioning the graph into a subset S and its complementary \bar{S} , such that the sum of weights in the set of cut edges is minimized relative to the total weight of internal edges in S . The sum of weights in the set of cut edges is analogous to a boundary term, while the total weight of internal edges is analogous to a volume term. In that sense the problem is an isoperimetric optimization problem, and the optimal set S minimizes the *isoperimetric ratio* ρ over all the possible subgraphs:

$$S = \operatorname{argmin}_{\operatorname{Vol}(s) \leq \operatorname{Vol}(\bar{s})} \rho(s), \quad (5.3)$$

$$\rho(s) := \frac{|\partial s|}{\operatorname{Vol}(s)}, \quad (5.4)$$

$$|\partial s| = \sum_{i \in s, j \in \bar{s}} W_{ij}, \quad (5.5)$$

$$\operatorname{Vol}(s) = \sum_{i \in s, j \in s} W_{ij}, \quad (5.6)$$

where the quantity W_{ij} is the weight of the link between vertex i and vertex j . Once S has been worked out, the algorithm can be applied separately to S and its complementary \bar{S} . This recursive application goes on until the isoperimetric ratio reaches a stopping ratio, representing the highest allowed isoperimetric ratio. This value is a parameter of the algorithm. Rising it results in a higher number of clusters, as it rises the number of acceptable cuts. The implementation of the recursive partition of the bipartite graph in the present toolbox is due to Grady and Schwartz [8].

Example 30. Partition the left hemisphere of the brain according to the expression of a given set of genes, chosen for their high divergence from a uniform expression. The code snippet below should reproduce Figures ??, among other things.

```

1 indsBrain = 1 : numel( brainFilter );
2 ann = cor.Annotations;
% consider the row indices in the data matrix constructed from
% voxels in the left hemisphere ('big12' annotation)
5 leftFilter = get_voxel_filter( cor, ann.filter( 5 ) );
6 indsBrain = make_volume_from_labels( indsBrain, brainFilter );
7 indsLeft = indsBrain( leftFilter );

```

5.2. CONSTRUCTION OF A BIPARTITE GRAPH FROM A VOXEL-BY-GENE MATRIX 51

```

8% normalise the sum of each column of the
9% data matrix so that each corresponds corresponds to a
10% probability density in voxel space
11 DNorm = normalise_integral( D );
12% rank the genes by descending order of localization
13 kldivergence = KL_divergence_from_uniform( DNorm );
14 [ vals, inds ] = sort( kldivergence, 'descend' );
15 plot_intensity_projections( make_volume_from_labels( D( :, inds( 1 ) ), brainFilter ) );
16% keep the numGenesKept most localized genes
17 numGenesKept = 150;
18 numGenesTot = size( D, 2 );
19 indsTaken = inds( 1 : numGenesKept );
20 voxelByGeneMatrix = D( indsLeft, indsTaken );
21% for speed and memory, keep the most intense
22% voxels construct a bipartite graph (the thresholdRank voxels
23% with highest total energy kept)
24 voxelIntensities = sum( voxelByGeneMatrix, 2 );
25 [ valsInd, indsInt ] = sort( voxelIntensities );
26 thresholdRank = 6351;
27 voxelIndsKept = indsInt( 1 : thresholdRank )
28 adjacencyMatrix = biclustering_graph_voxel_by_gene( voxelByGeneMatrix, voxelIndsKept );
29% consistency check for the size of the adjacency matrix
30 numInGraph = size( adjacencyMatrix, 1 );
31 numVoxKept = numel( voxelIndsKept );
32 numVoxKeptComputed = numInGraph - numGenesKept;
33 if numVoxKept ~= numVoxKeptComputed
34 display( [ numVoxKept, numVoxKeptComputed ] );
35 error( 'size of adjacency matrix inconsistent with data' );
36 end
37% choose a range of stopping criteria (numbers between 0 and 1:
38% maximum fraction of the weight of links that allowed to be cut
39% before the algorithm terminates
40 stoppingCrits = 0.2 : 0.02 : 0.3;;
41 for ss = 1 : numel( stoppingCrits )
42 [ voxels, genes, cluster ] = bicluster_graph( adjMat, 'iso', stoppingCrits( ss ),...
43 numVoxKeptComputed );
44 myClusters. voxels{ ss } = voxels;
45 myClusters. genes{ ss } = genes;
46 myClusters. cluster{ ss } = cluster;
47 myClusters. stoppingCrit{ ss } = stoppingCrits( ss );
48 % the number of clusters should be a growing function
49 % of the stopping criterion
50 numClust( ss ) = numel( cluster );
51 end
52% visualize the sums of gene-expressions
53% in each of the clusters
54 for ss = 1 : numel( stoppingCrits )
55 stoppingCriterionUsed = stoppingCrits( ss );
56 display( stoppingCriterionUsed )
57 cluster = myClusters. cluster{ ss };
58 for cc = 1 : numel( cluster )

```

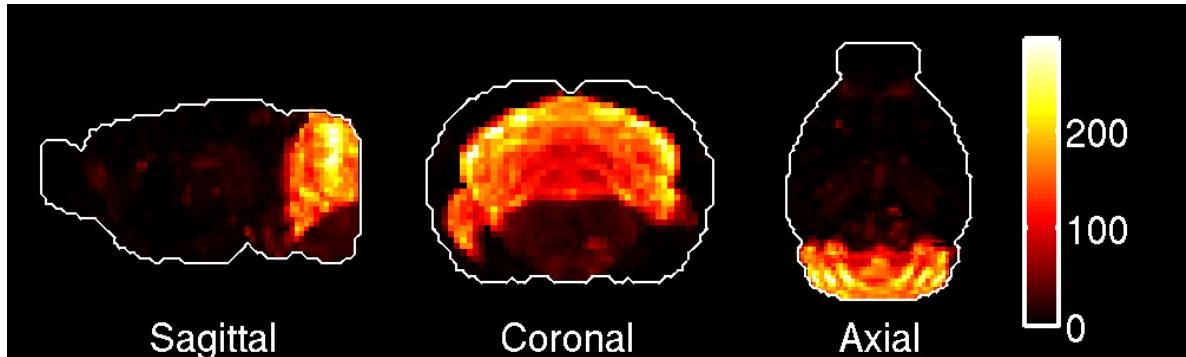


Figure 5.1: Maximal-intensity projections of the brain-wide sum of gene-expression for the genes in the first bicluster returned by the Example 31.

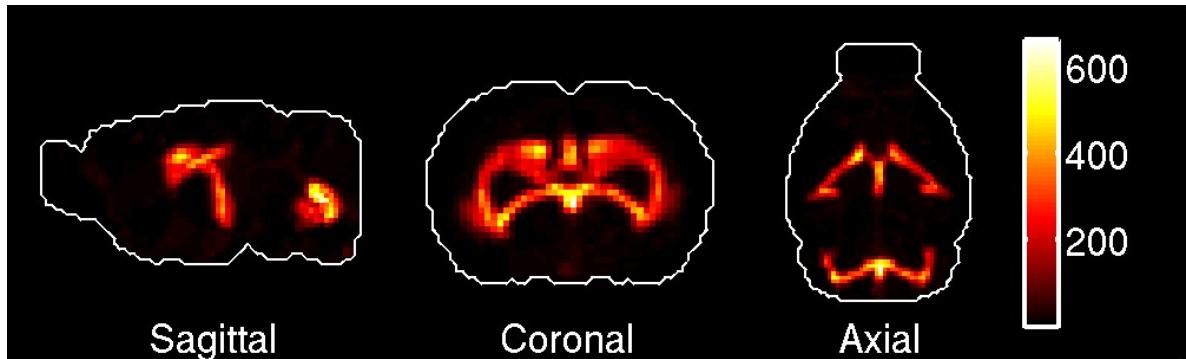


Figure 5.2: Maximal-intensity projections of the brain-wide sum of gene-expression for the genes in the second bicluster returned by the Example 31.

```

59     geneIndsInTaken = cluster( cc ).geneIdx;
60     indsInAtlas = indsTaken( geneIndsInTaken );
61     geneSum = sum( D( :, indsInAtlas ), 2 );
62     plot_intensity_projections( make_volume_from_labels( geneSum, ...
63         brainFilter ) );
64 end
65 pause;
66 hold off;
67 close all;
68end

```

5.2. CONSTRUCTION OF A BIPARTITE GRAPH FROM A VOXEL-BY-GENE MATRIX 53

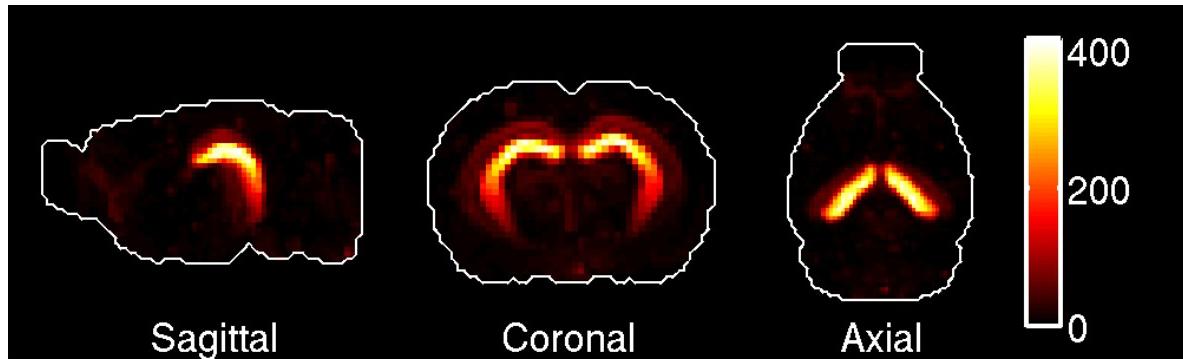


Figure 5.3: Maximal-intensity projections of the brain-wide sum of gene-expression for the genes in the third bicluster returned by the Example 31.

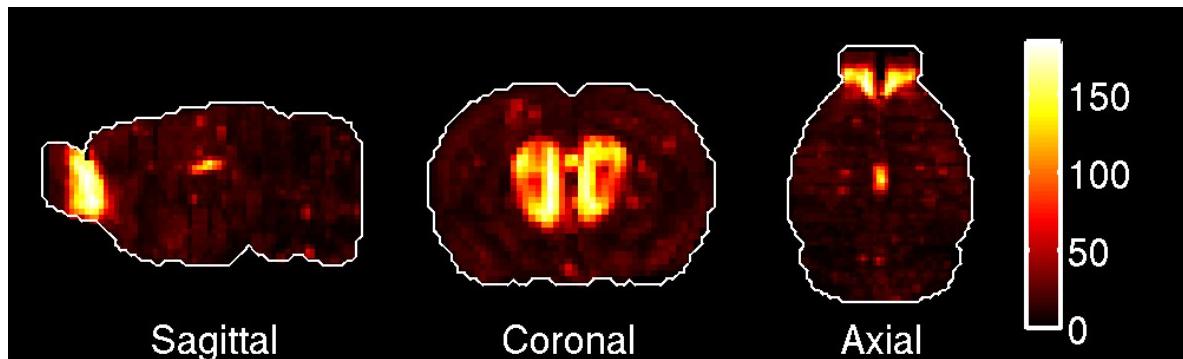


Figure 5.4: Maximal-intensity projections of the brain-wide sum of gene-expression for the genes in the first bicluster returned by the Example 31.

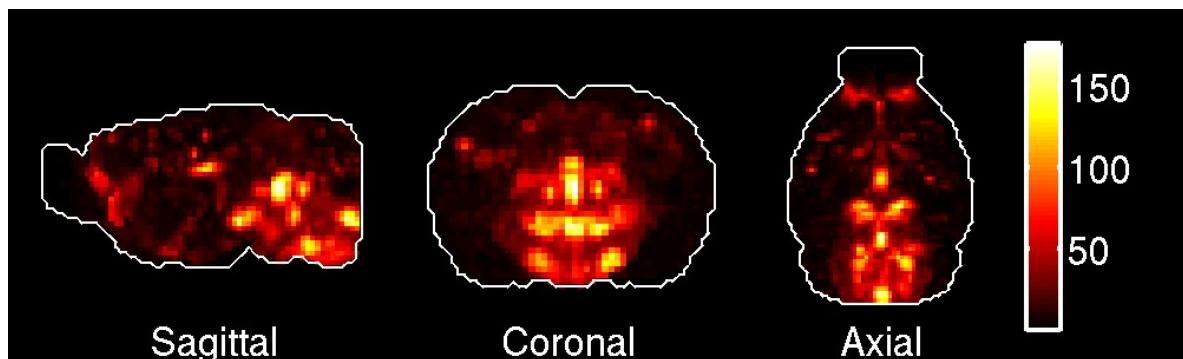


Figure 5.5: Maximal-intensity projections of the brain-wide sum of gene-expression for the genes in the second bicluster returned by the Example 31.

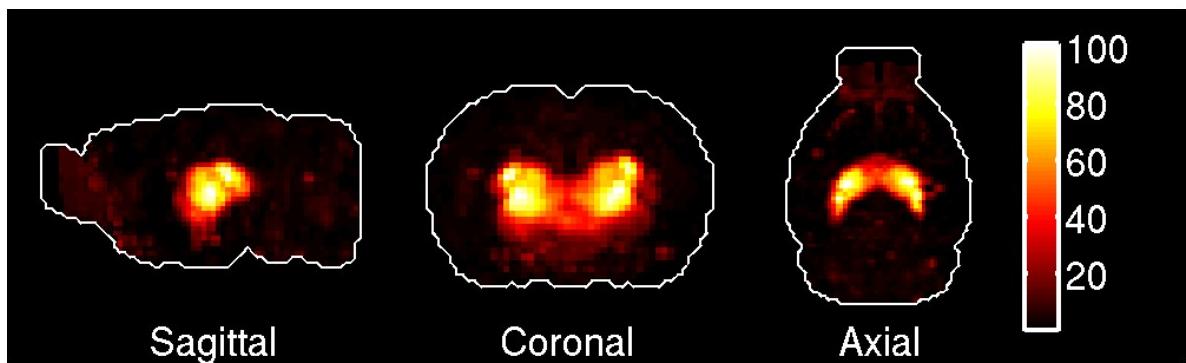


Figure 5.6: Maximal-intensity projections of the brain-wide sum of gene-expression for the genes in the third bicluster returned by the Example 31.

Chapter 6

List of files

6.1 Preparation of data

```
mouse_start_up.m  
cell_type_start_up.m
```

6.2 Atlas navigation

```
get_annotation.m  
get_genes.m  
get_voxel_filter.m  
make_volume_from_labels.m  
plot_intensity_projections.m  
flip_through_sections.m  
outline_section_with_atlas_picked.m  
annotation_big12_to_fine  
annotation_fine_to_big12.m
```

6.3 Marker genes

```
localization_from_id.m  
localization_scores_region_by_gene.m  
fitting_from_id.m  
fitting_scores_region_by_gene.m  
find_find_lambdamax_11_ls_nonneg.m (author K. Koh)  
11_ls.m (author K. Koh)
```

11_ls_nonneg.m (author K. Koh)

6.4 Co-expression analysis

```
co_expression_matrix.m
upper_diagonal_coeffs.m
cumul_co_expr.m
cumul_co_expr_plot.m
co_expression_components_bootstrap.m
co_expression_components_bootstrap_plot.m
```

6.5 Clustering

```
KL_divergence_from_uniform.m
biclustering_graph_voxel_by_gene.m bicluster_graph.m
partitiongraph.m (authors: A. Pothen, H. Simon and K.-P. Liou [9])
recursivepartition.m (author: L. Grady)
isosolve.m (author: L. Grady)
```

6.6 Cell-type specific analysis

```
cell_type_start_up.m
cell_types_correls.m
fit_voxels_to_types_picked.m
cell_type_vol_prepare.m
classify_pattern.m
figure_for_types.m
figure_for_type_correlations.m
figure_for_type_correlationsdensity.m
plot_correlations_for_big12.m
plot_densities_for_big12.m
```

6.7 Miscellaneous operations on matrices

```
normalise_integral.m.m
normalise_integral_L2.m
```

6.8 Data files

6.8.1 Allen Brain Atlas, Allen Reference Atlas

refDataStruct.mat
ExpEnergy.mat
ExpEnergytop75percent.mat

6.8.2 Cell-type specific analysis

G_t_means.txt
colsToUseInAllen.mat
colsToUseInTypes.mat
cellTypesDescription.mat

6.9 Saved results

6.9.1 Marker genes

quadTotIdentifier5.mat
quadRegIdentifier5Region1.mat
quadRegIdentifier5Region2.mat
quadRegIdentifier5Region3.mat
quadRegIdentifier5Region4.mat
quadRegIdentifier5Region5.mat
quadRegIdentifier5Region6.mat
quadRegIdentifier5Region8.mat
quadRegIdentifier5Region9.mat
quadRegIdentifier5Region10.mat
quadRegIdentifier5Region11.mat
quadRegIdentifier5Region12.mat
quadRegIdentifier5Region13.mat

6.9.2 Cell-type specific analysis

cellTypesCorrelations.mat
fitVoxelsToTypes.mat

6.9.3 Playground

`tool_box_manual_illustration.m`

Bibliography

- [1] H.-W. Dong, *The Allen reference atlas: a digital brain atlas of the C57BL/6J male mouse*, Wiley, 2007.
- [2] E.S. Lein, M. Hawrylycz, N. Ao, M. Ayres, A. Bensinger, A. Bernard, A.F. Boe, M.S. Boguski, K.S. Brockway, E.J. Byrnes, L. Chen, L. Chen, T.M. Chen, M.C. Chin, J. Chong, B.E. Crook, A. Czaplinska, C.N. Dang, S. Datta, N.R. Dee, *et al.*, *Genome-wide atlas of gene expression in the adult mouse brain*. *Nature* **445**, 168–176 (2007).
- [3] P. Grange and P.P. Mitra, *Computational neuroanatomy and gene expression: optimal sets of marker genes for brain regions*, IEEE, in CISS 2012, 46th annual conference on Information Science and Systems (Princeton).
- [4] K. Koh, S.-J. Kim and S. Boyd, *An interior-point method for large-scale L^1 -regularized logistic regression*, *Journal of Machine Learning Research* **8** (2007) 1519–1555.
- [5] P. Grange and P.P. Mitra, *Statistical analysis of co-expression properties of sets of genes in the mouse brain*, [[arXiv:1111.6200 \[q-bio.QM\]](https://arxiv.org/abs/1111.6200)].
- [6] S.F. Saccone, N.L. Saccone, G.E. Swan, P.A.F. Madden, A.M. Goate, J.P. Rice and L.J. Bierut, *Systematic biological prioritization after a genome-wide association study: an application to nicotine dependence*, *Bioinformatics* (2008), **24**, 1805–1811.
- [7] B.W. Okaty, K. Sugino and S.B. Nelson, *A Quantitative Comparison of Cell-Type-Specific Microarray Gene Expression Profiling Methods in the Mouse Brain*, *PLoS One* (2011) **6(1)**.
- [8] L. Grady, E. Schwartz, *Isoperimetric partitioning: A new algorithm for graph partitioning*, *SIAM Journal of Scientific Computing* **27** (2006), no. 6, 1844–1866.
- [9] A. Pothen, H. Simon and K.-P. Liou, *Partitioning Sparse Matrices with Eigenvectors of Graphs* (1990), **11**, 3, 430–452.